

**THE COMPLETE
TECHNICAL PAPER PROCEEDINGS**
FROM:



Published by:  **ncta** 

Compiled by:
Mark Bell, VP, Industry and Association Affairs
Wyatt Barnett, Senior Director, Industry and Association Affairs
Katie Mercier, Director, Programs & Events

Current and past editions of the *Technical Forum Proceedings* and *NCTA & SCTE·ISBE Technical Papers* are available online at www.nctatechnicalpapers.com.

ISBN Number: 0-940272-59-8
©2020, NCTA – The Internet and Television Association.
All rights reserved.

Steven Epstein - Synamedia
***Credential Fraud Detection And Remediation
In Media Consumption Services..... 1***

David Geary - Comcast
Robert F. Cruickshank III, Ph.D. - CableLabs
Derek DiGiacomo - SCTE
Ken Gilbert; John W. Teague; Mark Welsko, P.E. -
WES.net
Mike Glaser - Cox Communications
Brian T. Patterson - Emerge Alliance
Powering 10G: The Role Of Microgrids 12

Jason Rupe; Jingjie Zhu - CableLabs
***Profile Management Informed Proactive
Network Maintenance 35***

Shlomo Ovadia, Ph.D.; Deependra Rawat; Dan
Lynch - Charter Communications
***Streaming Telemetry Data from the Home
Network using OpenWrt Access CPE..... 63***

Ayham Al-Banna, Ph.D.; Tom Cloonan, Ph.D. -
CommScope
***DOCSIS 4.0 Network Migration Made Easy
..... 81***

Massimiliano Pala - CableLabs
***DOCSIS® PKI: A Proposal for a Next-
Generation Quantum-Resistant Infrastructure
..... 104***

Chris Seeger - NBC Universal, LLC
***Live Single-Stream HDR Production for HDR
and SDR Cable Distribution 127***

Karthik Krishna; Jambi Ganbar - VMware
***Reducing The Tradeoff Between Performance
And Management Using Container And Cloud-
Native Approaches 149***

Adrian Beaudin; Bruce Van Nice - Akamai
Building a Business Service in the Cloud 164

Mark Dokter; Bruce Van Nice - Akamai
***DNS Encryption: Exposure or Opportunity?
..... 172***

Hongbiao Zhang; Peter Wolff - Casa Systems
Dynamic IUC for OFDMA Transmission.... 186

Ron Hranac - Cisco Systems
Chad Campbell - Intraway
Roger Fish; Tom Kolze - Broadcom
Even Kristofferson; Aleksander Soeberg - Telia
Norge
James Medlock - Akleza
Jason Rupe; Tom Williams - CableLabs
Paul Schauer; Larry Wolcott - Comcast
Full Band Capture Revisited 198

Maher Harb; Bryan Santangelo; Dan Rice; Jude
Ferreira - Comcast
Full Scale Deployment of PMA..... 244

Nick Pinckernell; Scott Rome - Comcast
***Operationalizing Streaming Telemetry and
Machine Learning Model Serving: Customer
Experience Automation 267***

J.R. Flesch; Bryan Pavlich; Kurt Lumbatis; Charles
Cheevers - Commscope
***Thank You FCC -- Now We Have 1.2 Ghz of 6
GHz Spectrum so How Does the Cable
Operator Utilize It? 290***

John J. Downey - Cisco Systems
***The Power of Distributed Access Architectures
(DAA) 316***

John J. Downey - Cisco Systems
***Guide to Managing Cable Network Traffic
Congestion..... 331***

Ian Wheelock; Charles Cheevers - CommScope
Dr. Sudheer Dharanikota; Ayarah Dharanikota -
Duke Tech Solutions Inc.
***The Business Case for Aging in Place with
Cable Operators 344***

Dr. Sudheer Dharanikota; Ayarah Dharanikota -
Duke Tech Solutions Inc.
***Why Are Cable Operators A Natural Fit To
Support Telehealth: An Inter-Industry
Perspective..... 383***

Rafie Shamsaasef; Aaron Anderson; Sasha
Medvinsky - CommScope
***Cloud-based Dynamic Executable Verification
..... 395***

Timothy Maenpaa - Ciena Corporation
***Addressing Unrelenting Growth In Backbone
Fiber Systems Using Next Generation
Photonics And Automation..... 420***

Tom Woginrich - International Business Machines
(IBM)
***IoT & Cognitive Computing Approach to
Managing Equipment 434***

Helen Zeng, Ph.D.; Robert McIntyre - VMware
***Employing Neural Networks For Improved
Root Cause Analysis In Service Provider
Clouds 446***

Ladan Pickering - Fujitsu Network Communications
***City Of Dublin: Lessons From A Smart City
Private Network Deployment 462***

Jason Page - Charter Communications
***A Better Platform to Facilitate Remote Patient
Monitoring..... 471***

Qi Zhou; You-Wei Chen; Shuyi Shen; Gee-Kung
Chang - Georgia Institute of Technology
Jeff Finkelstein; Drew Davis; Brian Lee - Cox
Communications
***Simultaneous Echo Cancellation and
Upstream Signal Recovery using Deep
Learning in Full-duplex DOCSIS Systems..484***

Eric Heaton - Intel Corporation – Network Platforms
Group
***Strategies for Implementing Edge Services in
the 10G Cable Network 494***

Jeff Finkelstein - Cox Communications
Tom Cloonan - CommScope
Doug Jones - CableLabs®
***Covid-19 Learnings: All Roads Lead To
DOCSIS® 4.0 Technology..... 518***

Alexander Medvinsky; Dr. Tat Chan; Dr. Xin Qiu;
Jason Pasion - CommScope
***A Flexible and Scalable Architecture for Over-
the-Air Credentials Provisioning 537***

Wesley Weiss; Anjan Bajwa; Corwin Martens - Shaw
Communications Inc.
***Leveraging Legacy Video In Digital Access
Architecture Networks 562***

Kyle Hohman - Shaw Communications
***To High Split & Beyond: The New Frontier In
Leakage Detection 592***

Colin Dearborn - Shaw Cablesystems G.P.
***Improving The Latency Of An MSO Network
For Gaming And Real Time Applications ... 605***

Nader Foroughi - Shaw Communications
Jason Rupe - CableLabs®
***Distributed Gain Architecture: Increased
Performance, Decreased Power Draw 624***

Dr. Thushara Hewavithana - Intel Corporation
Dr. Rainer Strobel
Nader Foroughi - Shaw Communications Inc
***Closed Loop Capacity Optimization for
Extended Spectrum DOCSIS..... 674***

Patricio Sebastian Latini - Casa Systems
***The Headend Evolution: Design
Considerations for Deploying vCCAP and
Other VNFs 704***

Todd Musat - Shaw Communications Inc.
Chuck Carroll; Lew Rakowsky; Rene Spee - Saras
Energy Consulting
***Powering The Future: Next Generation Access
Networks..... 729***

Chris Day; Joshua Rose - Analog Devices
***Optimizing Active Components for Extended
Spectrum Networks 750***

Kevin Taylor; Michael Khalilian - Comcast
Steve Goeringer - CableLabs
Eric Winter - Cox Communications
***A Taxonomy of Fraud Experienced by Network
Service Providers 770***

Will Bracker - Cox Communications
Steve Goeringer; Simon Krauss - CableLabs
***Fraud Prevention and Privacy Law: Emerging
Conflicts Between Privacy Law and Fraud
Prevention 780***

Sebnem Ozer, Ph.D.; Carl Klatsky; Dan Rice; John
Chrostowski - Comcast
***Approaches to Latency Management:
Combining Hop-by-Hop and End-to-End
Networking 788***

Anthony Curran; Andy Martushev - Comcast
***Software Revolution Of Field Meters Using a
Field-Capable Measurement Device811***

Tong Liu, PhD; John T Chapman - Cisco Systems
Inc
***Building the RPHY Upstream Scheduler with
YANG.....827***

Gary Ventriglia; Jack Birnbaum; Robert Gonsalves;
Anastasia Vishnyakova; Michael Kreisel; Larry
Wolcott - Comcast Corporation
***Training Machines to Learn From Signal
Meter Readings: A Case Study from Comcast
.....846***

Sameh Yamany; Paul Gowans - VIAVI Solutions
***Exposing The Invisible Enemy: How Network
Location Intelligence and Analytics Saves
Lives.....871***

Srilal Weerasinghe PhD - Charter Communications
***MLaaS Applications in Digital Video –
Supplanting Disliked Content.....878***

Yoshitaka Kidani; Hiroyuki Yamashita; Shuichi
Matsumoto - Japanese Cable Laboratories
***Proposal of RF/IP Adaptive Video Distribution
Scheme over Cable Television Access Networks
.....891***

Tom DiMicelli - Ciena Corporation
***Unleashing Managed SD-WAN With Closed-
Loop Automation.....902***

Venk Mutalik; Bob Gaydos; Dan Rice; Jorge
Salinger - Comcast
***Accelerating the Virtualization: Introducing
Hybrid Fiber Shelf into the Mix918***

Venk Mutalik; Dan Rice; Bob Gaydos; Doug Combs;
Pat Wike - Comcast
***Operationalizing the Grey Optics Architecture:
An Update One Year After.....941***

Venk Mutalik; Dan Rice; Rick Spanbauer; Simone
Capuano; Rob Gonsalves; Bob Gaydos - Comcast
***It's 10 PM: Do You Know Where Your
Wavelengths Are?966***

William McFarland - Plume
***Connectivity and COVID-19: Maintaining QoE
During a Crisis990***

Frank Sandoval - Pajarito Technologies, LLC
***Power Management on the Generic Access
Platform.....1015***

Fernando X. Villarruel; David Reale - Ciena
***Framework for Convergence of Services on
The MSO Network1029***

Stuart Kurkowski, PhD; Neill Kipp - Comcast
Technology Solutions
***Satellite to Fiber Broadcast Execution With
SCTE 35 and 224.....1049***

Robert Cruickshank, Ph.D. - Cable Television
Laboratories / National Renewable Energy
Laboratory
Nicolas Metts - Cable Television Laboratories
Paul Schauer - Comcast Cable Communications
Curtis Snyder - Consultant
***Gridmetrics™ Data Provide Insights and
Improve Situational Awareness of the Electric
Power Grid1073***

Rob Anderson - EnerSys Energy Systems
***Powering the Near Future 10G Access
Network: Considerations for Assuring
Sufficient and Reliable Power.....1089***

Javier Ger - Telecom Argentina
Claudio Saes - Bell Labs Consulting
***From CSP to DSP: Is the COVID-19 Crisis a
Partner or Another Steppingstone?1116***

Karthik Sundaresan; Jay Zhu - CableLabs
João Pedro Fernandes - NOS
***Field Experiences with US OFDMA and using
US Profile Management1144***

Karthik Sundaresan; Greg White; Steve Glennon -
CableLabs
***Latency Measurement: What is Latency and
How Do We Measure It?1172***

Kyle Haefner - CableLabs
***Smart Gateways: Active A.I. in Subscriber
Networks.....1197***

Elías Chavarría Reyes, Ph.D.; John T. Chapman -
Cisco Systems, Inc.
***How DOCSIS Time Protocol makes the SYNC
Specification Tick1212***

John Chrostowski; Greg Tresness; Dan Rice; Benny
Lewandowski – Comcast
Greg Tresness – Arcom Digital, LLC.
***Leakage In A High Split World* 1233**

Thomas Hurley - Comcast Cable Corporation
John Dolan - Rogers Communications Inc.
Arnold Murphy - Strategic Clean Technology Inc.
Mike Glaser - Cox Communications Inc.
John Teague - Worldwide Environmental Services
Ken Nickel - Quest Controls
***Critical Facility Cooling Energy Optimization*
..... 1271**

Rob Thompson - Comcast
Xiaohua Li - State University of New York at
Binghamton
***Machine Learning Techniques for Equalizing
Nonlinear Distortion* 1292**

Andrii Vladyka; Asaf Matatyau - Harmonic Inc.
***Virtualization and Edge Compute Evolution in
Cable* 1330**

Venkata Somi Alapati; Kashif Shakil - Ericsson
Craig Schwechel - Incode Consulting
***Path To Gigabit Fixed Wireless Access: A
Review Of Fixed Wireless Access Technology
And Economics* 1342**

Dr. Robert Howald; Robert Thompson; Sebnem
Ozer; Daniel Rice; Larry Wolcott - Comcast
Dr. Tom Cloonan; Dr. Ruth Cloonan; John Ulm -
CommScope
Jan Ariesen - Technetix
***Roaring Into The '20s With 10G* 1370**

Dr Vikram Saxena; Ryan Eccles - NetScout Systems
***Smart Data Powers Service Layer Management
For Network Operations 2.0* 1418**

Matthew Schmitt - CableLabs
***Constructing a Convergence Lab* 1427**

Keith R. Hayes - IMMCO, Inc.
***Using Machine Learning To Automate Node
Split Designs And HFC Augmentation Options*
..... 1448**

Bill Beesley - Fujitsu
***Bringing Service Visibility Into The Light With
CPRI As A Service* 1461**

Bill Beesley - Fujitsu
***Make the Most of What You've Got: How
Cable Modems Can Deliver Economical Cell
Site Transport* 1469**

Charles Cheevers; Ian Wheelock; Kurt Lumbatis -
CommScope
Kamal Koshy; Ahmed Bencheikh - Charter
Communications
***With 1.2GHz of Spectrum Are We Moving to a
Channelized per Room Architecture for the
Home – Enabled by Wi-Fi 7* 1475**

Ethan Wright - Charter Communications
***Augmented Reality Can Improve Wi-Fi
Installation In Homes* 1513**

Wael Guibene; Hossam Hmimy - Charter
Communications Inc
***Enabling Automatic Gunshot Detection and
First Responders Dispatch for Safer
Communities* 1525**

Dr. Robert Howald - Comcast
***Repair The Ides Of March: COVID-19
Induced Adaption of Access Network Strategies*
..... 1535**

Craig Pratt; Darshak Thakore - CableLabs
Jacob Gladish - Comcast
***Wi-Fi Passwords: The Evolving Battle Between
Usability and Security* 1568**

Wael Guibene; Hossam Hmimy - Charter
Communications
***Enforcing Social Distancing Using Computer
Vision and Deep Learning* 1590**

Bruce E. Barker Jr. - Next Generation Access
Network (NGAN), Comcast Cable
Claude Bou Abboud; Erik Neeld - Comcast Cable
***Access Capacity Planning: Staying Well Ahead
Of Customer Demand Helped Ensure Stability
During COVID-19* 1601**

John T Chapman - Cisco Systems
***Small Cell Traffic Engineering* 1618**

Srinath V Ramaswamy - Comcast
***Content Aware Video Streaming* 1644**

Jude Ferreira; Maher Harb; Karthik Subramanya;
Bryan Santangelo; Dan Rice - Comcast
***Convolutional Neural Networks for Proactive
Network Management*..... 1654**

Jorge Salinger; Steve Sigman - Comcast Cable
Communications
***DAA Field Deployment, Path to Scaling, and
Digital Node Use Cases*..... 1669**

Umamaheswar Achari Kakinada; Dr. Hossam H.
Hmimy; Manish Jindal - Charter Communications
Inc.
***A Method and Framework for IoT Network
Security*..... 1697**

Ram Ranganathan; Chris Markovic; Tushar Mathur;
Omar Abu-Hijleh; Thomas Cloonan; John Ulm -
CommScope
***Decoding The Bandwidth Surge During Covid-
19 Pandemic*..... 1709**

Alexander Giladi - Comcast Cable
***Session Overhead Reduction in Adaptive
Streaming*..... 1726**

Jennifer Andréoli-Fang, PhD - CableLabs
John T. Chapman - Cisco
***Cable and Mobile Convergence*..... 1739**

Parmjit Dhillon; Mohamed Daoud - Charter
Communications Inc
***Collecting Smart City IoT Data to Generate
Actionable Insights*..... 1827**

Richard S Prodan, Ph.D. - Comcast
***Optimizing the 10G Transition to Full-Duplex
DOCSIS® 4.0*..... 1845**

John Jason Brzozowski - MachineQ, a Comcast
company
***Bringing Enterprise IoT to Cable*..... 1881**

Dr. Mehmet Yavuz - Celona, Inc.
***Private Mobile Networks - A New Service
Option for Enterprise Wireless Connectivity*
..... 1891**

Marcus Rebelo - Resolve Systems
***In Pursuit of the Dark NOC: Driving Change
With Automation & AIOps*..... 1899**

John Ulm; Dr. Zoran Maricevic; Dr. Frank O’Keeffe
- CommScope
Is “Unity Gain” Still the #1 Objective? 1909

Dr. Claudio Righetti; Mariela Fiorenzo; Omar
Hurtado; Gabriel Carro - Telecom Argentina S.A.
***Augmented Intelligence: Next Level Network
and Services Intelligence*..... 1949**

John Ulm; Dr. Thomas Cloonan - CommScope
***Managing the Coronavirus Bandwidth Surge:
How to Cope with the Spikes and Long-term
Growth*..... 1979**

Tushar Mathur; Ram Ranganathan; Greg Gohman -
CommScope
Bob Zhang - University of Waterloo
***Low Latency Docsis: Concepts and
Experiments*..... 2008**

Jay Bestermann - CommScope
Ken Florenz - Altice USA
***Enterprise Opportunities Beyond The Pipe:
The Second Network*..... 2025**

Louis Donofrio; Qin Zang - Comcast Cable
Vignesh Ramamurthy - Infosys Consulting
***Key Learnings from Comcast’s Use of Open
Source Software in the Access Network*..... 2037**

Jason Combs - Comcast
***A Virtual Broadband Network Gateway
(vBNG) Approach for Cable Operators in a
Distributed Access Environment*..... 2046**

Mark Francisco - Comcast Cable
***The Future of Cable Television Audio is
Accessible*..... 2057**

Lei Zhou; Robert Thompson; Robert Howald; John
Chrostowski; Daniel Rice - Comcast Cable
***A Proactive Network Management Scheme for
Mid-split Deployment*..... 2068**

Andrew Frederick - Comcast
***Case Study of Social Distancing on Mentorship
Programs: CLEAR Program Introduction*
..... 2081**

Andrew Frederick - Comcast
***Expediting New Product Deployments with
Agile Operations and DevOps*..... 2095**

Robert M. Lund - Comcast
Kathryn Sanders – Sanders RF Consulting LLC
***RF Testing Applications for Software-Defined
Radio..... 2108***

Elliott Hoole - Charter Communications Inc
***Enabling Industry 4.0 Business Models For
MSOs Using Wireless Mesh Networks At 60
GHz..... 2133***

Paul E. Schauer - Comcast Cable
50 Million Keys to SNMPv3 Privacy 2160

Rohini Vugumudi; Hany Fame; Pardeep Singh; Zhen
Lu - Comcast
***Dynamic Data Collection & Configuration
Management..... 2172***

Joshua Seiden; Nishesh Shukla - Comcast Innovation
Labs
***Augmented Reality for Network Visualization
..... 2192***

Shawn Kercher; Jacob Hallberg - Comcast
Innovation Labs
***Verification of Electrical Grounds/Bonds
Using Computer Vision..... 2200***

Colleen Szymanik - Comcast
Aggregate Wi-Fi Telemetry Use Cases..... 2213

Albert Garcia - Comcast
***Developing Installation Guidelines For Wi-Fi
Managed Devices 2222***

Matthew Tooley; William A. Check, Ph.D.; Rob
Rubinovitz; Jim Partridge - NCTA – The Internet &
Television Association
***Tele-Everything and Its Impact to The Network
..... 2231***

Haider Syed - Charter Communications, Inc.
***Wireless Access Network Strategies: Lessons
Learned On 3.5 GHz CBRS Network Trials
..... 2254***

Massinissa Lalam, Ph.D; Xavier Briard - Sagemcom
Kamal Koshy - Charter Communications
***Wi-Fi 6 And Wi-Fi 6E Are Building The
Foundation For New Home Applications.. 2291***

Eli Baruch - CommScope
Indira Paudel - Comcast
***Using Big Data To Fine-Tune The Nation's
Largest Public Wi-Fi Network 2309***

Kevin Bourg; Sergey Ten; Peter Wigley - Corning
Optical Communications
***An Overview Of Optical Architectures
Necessary To Achieve 5G's Key Performance
Indicators 2331***

Marco Naveda; Dmitri Fedorov; Raghu Ranganathan
- Ciena
***Delivering Cloud-Native Operations with Edge
Compute Enabled DAA 2342***

Credential Fraud Detection And Remediation In Media Consumption Services

A Technical Paper prepared for SCTE•ISBE by

Steven Epstein

Distinguished Engineer
Synamedia

054-566-4116
sepstein@synamedia.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. The prevalence of fraudulent password usage in media services.....	4
3. Why credential fraud detection is so difficult: Sharing vs Fraud.....	6
4. Credential Fraud indicators and Solution Enablement.....	8
5. Trusted Identity as a Service.....	9
6. Conclusion.....	10
Abbreviations.....	11
Bibliography & References	11

List of Figures

Title	Page Number
Figure 1 - Credential Stuffing Attack.....	4
Figure 2 - Credential Stuffing Attack Rate by Industry	5
Figure 3 - Daily Malicious Login Attempts Against media	5
Figure 4 - Monthly Malicious Login Attempts Against Video Media	6
Figure 5 - Online Sales of Fraudulent Subscription Services	7
Figure 6 - Synamedia Anti-Fraud Algorithm.....	8
Figure 7 - Anti-Fraud Detection and Resulting Policy	9

List of Tables

Title	Page Number
Table 1 - Fraud Policy Table.....	10

1. Introduction

Accessing another's credentials has always been a major goal of hackers or pirates. Typically, pirates would perform phishing or even spear phishing attacks on naïve or unsuspecting targeted individuals. In these attacks, the hacker would send a user a link to some embellished website mimicking a known banking, credit card or other financial site. It would request the unsuspecting user to enter their personal credentials. Once the credentials were entered and transferred to the pirate, the pirate could now perform bank transfers, embezzle money or even take over the victim's account. These attacks were very costly to financial institutions and other highly secure websites, but not highly effective or scaleable. That's because in order to be successful, phishing sites required much intelligence to send the proper link to the appropriate users and even so most users did not take the bait.

In the last five years however, a new more scaleable and effective method of accessing another's credentials has become increasingly popular. This form of piracy, known as credential stuffing, is based on two historical realities:

1. In the past 10 years, thousands of identity databases belonging to large websites, have been breached leading to the identity theft of tens of billions of credentials.
2. Most people reuse the same credentials (username and password) on multiple sites as a convenient way of remembering them.

Given these two facts, new credential stuffing tools were created to enable a set of bots over proxies or VPNs to discover active breached credentials from a set of popular websites. The diagram below illustrates how credential stuffing attacks are performed.

A pirate purchases millions of username/password combinations (combos) extracted from breached websites, and configures a set of bots, proxies, desired websites and scripts describing login navigation details of each of these desired sites. The pirate then inputs all these artifacts into credential stuffing tools. The tool then assigns bots to try all of these combos on each of the popular websites, using navigation instructions within scripts, and connect to them via separate proxies or VPNs. In order to go undetected, different IP addresses are used for each malicious attempt! The tool returns a subset of the list of credentials that are still active on each popular site. This attack is effective because most users tend to employ the same username/password combination across most of their websites.

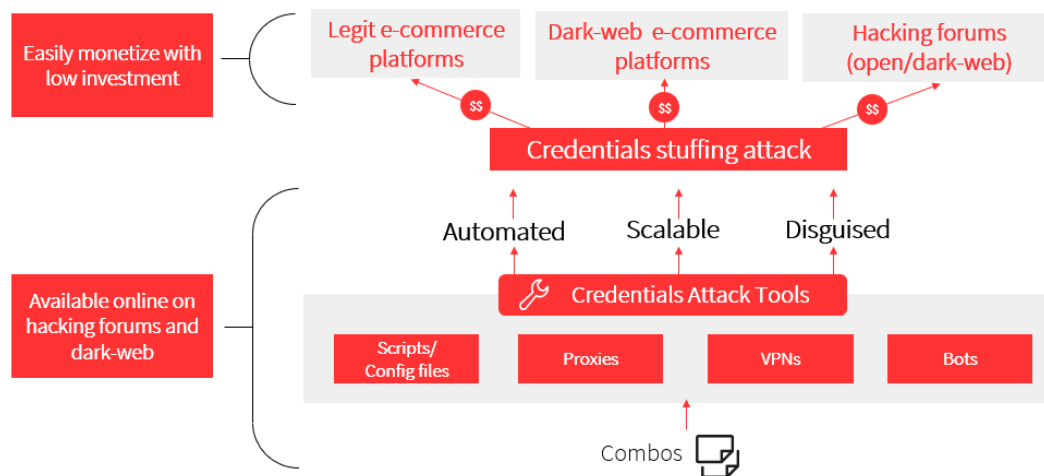


Figure 1 - Credential Stuffing Attack


Once any login attempt succeeds, the stolen credentials become marketable and can enable access to another's account. The accounts which are typically accessed via credential stuffing are not only banking sites but also streaming media sites, where the purpose of the attack is not to pilfer money, but to sell access to someone else's account and enable a buyer to enjoy content made available by another's monthly subscription. As a result of these credential stuffing attacks, access to media sites such as streaming video, music, audio book, books etc. are sold over telegram groups, forums, websites on the open Internet and especially the dark web.

Unlike phishing or spear phishing attacks, where the victim realizes the credential theft and resulting fraud fairly early and immediately alerts authorities, credential stuffing attacks on media sites tend to go undetected for a very long time. That's because the pirate or purchaser of the credentials desires to receive a free media service and hence does whatever it takes to remain unnoticed by the real owner.

2. The prevalence of fraudulent password usage in media services

Credential stuffing is growing dramatically, especially in media services. According to Forbes, 4.1 billion credentials were breached in the beginning of 2019 alone. The [HaveIBeenPwned](#) website has a database of over 10 billion breached credentials. According to Shape Security, between 0.5 and 2% of these credentials will be valid on any targeted website or mobile app.


Between the beginning of May and the end of June this year, Akamai collected data on 8.35 billion credential stuffing attempts across the globe across all industries. They estimate that in Q1, credential stuffing attacks increased by 1,450%, compared with about 200% in 2019 (compared with 2018). One European broadcaster was even hit with peaks of malicious credential stuffing attempts that ranged into the billions. However, as mentioned above, unlike phishing attacks which typically target financial institutions, a good percentage of credential stuffing attacks are focused on media. The graph below by Arkose Labs shows that over 20% of all online traffic to media sites are credential stuffing attacks.



Data source: Arkose Labs

Figure 2 - Credential Stuffing Attack Rate by Industry

These attacks against media sites continue to grow significantly. The graph below by Akamai exhibits the steep rise in malicious credential stuffing login attempts against media, exacerbated by the period of COVID-19, where a 300% increase was observed.



Daily malicious login attempts during Q1 2020. Source: Akamai

Figure 3 - Daily Malicious Login Attempts Against media

But even before COVID-19, between 2018 and 2019 there was a 63% increase in credential stuffing attacks against video media sites as shown by the Akamai graph below.




Figure 4 - Monthly Malicious Login Attempts Against Video Media

Akamai researchers watching the credential stuffing space in Q1 2020, noted that video media accounts were trading for about \$1 to \$5 on the criminal market early on. Some packaged offers (those that include multiple services per order) were even being sold for \$10 to \$45. Toward the end of Q1 2020, those prices fell as the credential stuffing market became flush with new accounts and lists of recycled credentials.

Based on the statistics above, it is clear that the criminal economy is a chained instance, where everything is somehow connected, and no piece of information is without worth. Criminals pre-package compromised accounts, selling them based on interest, location, and volume, and people are willing to pay. This only fuels the criminals' actions and keeps them hyper-focused on evading detection and mitigation.

3. Why credential fraud detection is so difficult: Sharing vs Fraud

As mentioned above, credential stuffing discovers marketable credentials that are active on various media sites. These credentials are then typically sold on either the dark web, telegram groups, forums or open Internet websites. Below is an example of a site selling various subscriptions to streaming services for one time payments under \$15.

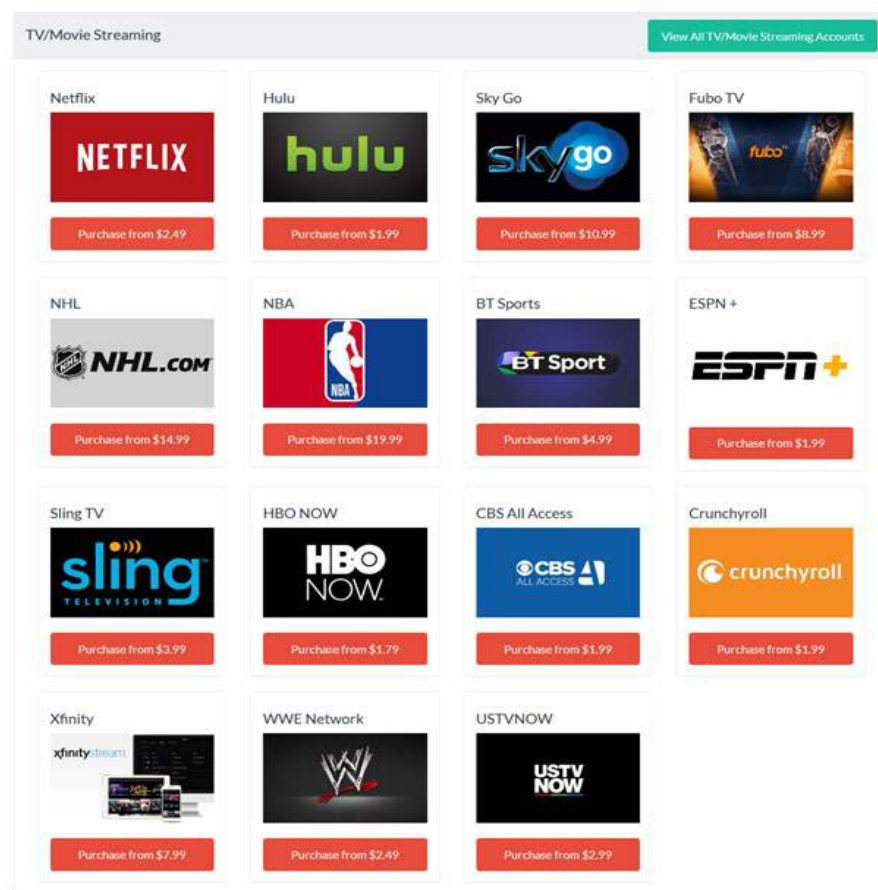


Figure 5 - Online Sales of Fraudulent Subscription Services

The motive of those purchasing these credentials is not malicious. They aren't looking to takeover and control the account. Nor are they looking to purchase goods on the account owner's credit card. They simply want to enjoy a popular streaming subscription service for a low one-time cost.

In other words, those who purchase another person's credentials to a popular video media site is unable to find another account owner who is willing to casually share their credentials with them. Hence they need to rely on purchasing the credentials from a stranger in order to receive a free subscription to this media service.

The goal of the user who purchases credentials from a credential stuffing fraudster is to enjoy this service that he purchased for a low one-time fee for as long as possible without being noticed and without the account owner changing his password out of suspicion. Hence, the behavior of the fraudster will by definition mimic that of the sharer, one who benefits from casual sharing, or even the account owner.

Based on this new reality, it becomes challenging by scrutinizing the data alone to differentiate between the account owner and the fraudster and even more difficult to differentiate between the fraudster and the casual sharer. If you scrutinize which devices within any given account constantly view video from out-of-home locations and IP addresses, you cannot distinguish between the fraudster, the person who purchased the fraudulent credentials or the casual sharer.

While most service providers are willing to tolerate some level of casual sharing, the same is not the case when it comes to credential fraud. Given concerns over privacy and other liabilities, detecting the fraudster and differentiating them from the sharer is critical to the well-being of a video service.

Which begs the question: How does one differentiate between casual credential sharing and credential fraud?

4. Credential Fraud indicators and Solution Enablement

Synamedia has created a novel solution to detect credential fraud in media services resulting from credential stuffing attacks based on machine learning. In this solution, credentials sold online to any particular media service are purchased by our own intelligence group. The credentials are then used to train a supervised machine learning model to detect a fraudulent account using various indicators.

Indicators of fraud in the model are based on some of the following principles:

1. Landscape of how credentials are sold on the various marketplaces and various trends on the frequency of usage of a stolen credential once put up for sale
2. Assumption that there is less correlation between the viewing preferences, behaviors and habits of the fraudster and the account owner versus those of the sharer and the account owner
3. Other anomalous, suspicious and unexpected activity in the account

Once this supervised model is built, it can be applied to locate other, as yet undiscovered fraudulent accounts. Once all fraudulent accounts are classified, a second model is built which can differentiate between the devices of the sharer and account owner and that of the fraudster within each classified fraudulent account. This algorithm is semi-supervised but is based on many of the same indicators described.

Finally, fraudulent devices are graded according to the likelihood that they will perform real malicious activity, such as an online purchase of goods and/or new services in the name the account owner or even an account takeover. The activity diagram describing this algorithm is shown below.

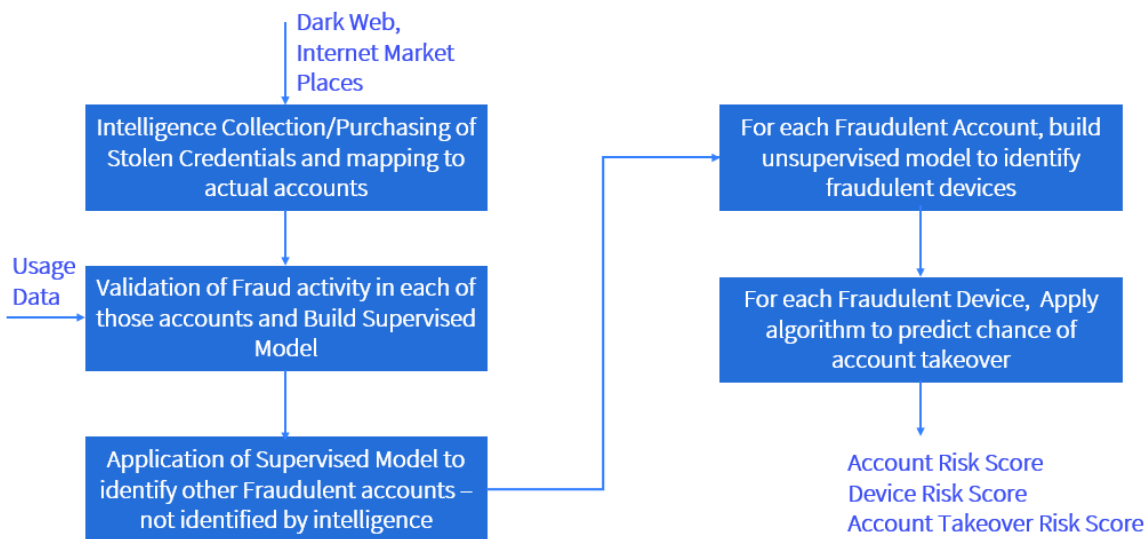


Figure 6 - Synamedia Anti-Fraud Algorithm

5. Trusted Identity as a Service

By creating such an algorithm, Synamedia provides a service to validate the trusted identity of media accounts and devices within each account. Upon receiving near real-time data streams from a media provider, Synamedia supports an API where it can return:

1. Likelihood that a specific account is fraudulent
2. Likelihood of fraud in each device within that fraudulent account
3. Risk of account takeover or other malicious activity of each fraudulent device in that fraudulent account

An example of this service is shown in the diagram below

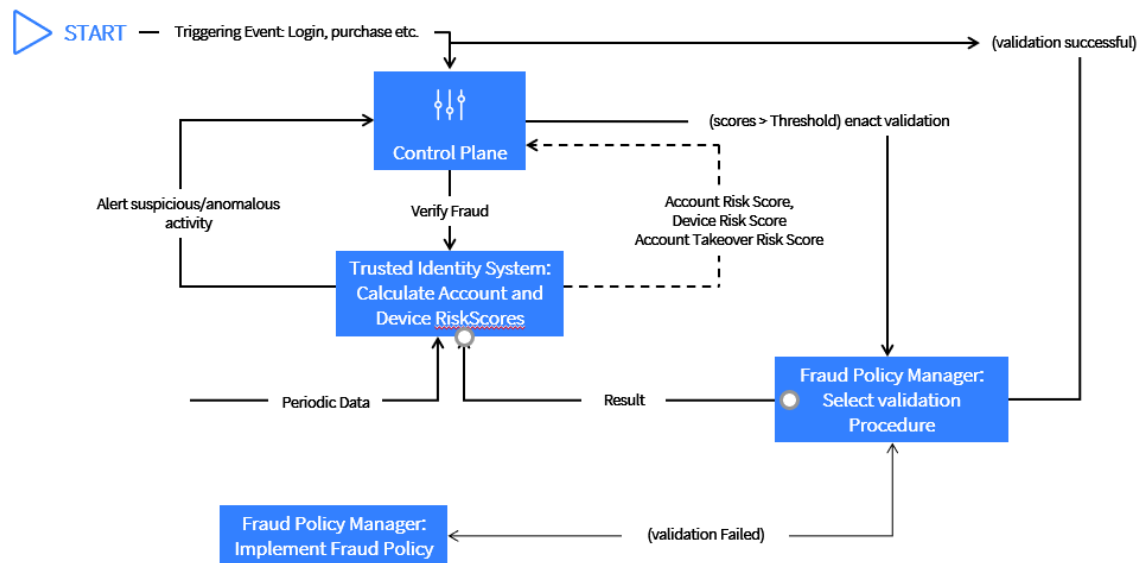


Figure 7 - Anti-Fraud Detection and Resulting Policy

This trusted identity service enables a media provider to both perform adaptive authentication and enforce adaptive remediation policies, where the authentication and the resulting policy are compliant to the risk of fraud of every device in every account. As a result, the service provider can create an adaptive policy table as a function of the sensitivity of the trigger (login, TVOD purchase, change of address etc.) which caused the trusted identity API to be called.

Many verification and remediation policies, such as password change and answering security questions, are effective against fraudsters as opposed to sharers because we assume no social connection nor passing of information between the account owner and the fraudster!

An example of an adaptive policy table is shown below.

Table 1 - Fraud Policy Table

Trigger Sensitivity Score	Account Risk Score	Device Risk Score	Account Takeover Risk Score	Fraud Verification Policy	Fraud Remediation Policy if verification Fails	Fraud Remediation Policy if verification Succeeds
<30	All	ALL	All	None	N/A	N/A
>30 <60	<50	<30	N/A	None	N/A	N/A
>30 <60	>50	<30	N/A	Provide Password	Suspend Account	Change Password
>30 <60	>50	>30	<20	Answer 1 or more Security Questions	Blacklist Device	Change Password
>30 <60	>50	>30	>20	Biometric / MFA	Suspend Account	Change Password
>60	<50	<30	N/A	Provide Password	Suspend Account	Change Password
>60	>50	<30	N/A	Answer 1 or more Security Questions	Blacklist Device	Change Password
>60	>50	>30	<20	Answer 1 or more Security Questions	Suspend Account	Change Password
>60	>50	>30	>20	None	Suspend Account and Blacklist Device	Change Password

6. Conclusion

Credential stuffing is a new piracy attack that has become prevalent in the past several years. This is due to the vast increase in the quantity of breached credentials in the marketplace and also based on the fact that few users change their credentials between accounts.

Credential stuffing, as opposed to other credential discovery techniques such as phishing, is prevalent not only on financial sites but also on media sites where the primary objective of the fraudster is to benefit from a low one-time payment to a subscription service. Hence, fraudulent use based on credential stuffing is extremely difficult to detect and disrupt!

Using advanced machine learning techniques, Synamedia has managed to build new fraud detection algorithms that can detect fraudulent usage of devices within accounts based on credential stuffing attacks. The results of this algorithm enable very effective adaptive authentication and remediation techniques to combat those passwords that were stolen and purchased based on credential stuffing!

Abbreviations

API	Application programmable Interface
CSA	Credential Stuffing Attack
TVOD	Transactional Video On Demand
VPN	Virtual Private Network
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

<https://informationisbeautiful.net/visualizations/worlds-biggest-data-breaches-hacks/>

https://piracymonitor.org/fraudulent-logins-q2-2020-arkose-labs/?utm_source=Piracy+Monitor&utm_campaign=3e0dea891d-PM-E-Newsletter-2019-1223_COPY_01&utm_medium=email&utm_term=0_baec17a8d9-3e0dea891d-364275657

<https://www.akamai.com/us/en/multimedia/documents/state-of-the-internet/soti-security-credential-stuffing-in-the-media-industry-report-2020.pdf>

<https://haveibeenpwned.com/>

Powering 10G: The Role Of Microgrids

A Technical Paper prepared for SCTE•ISBE by

David Geary, P.E., Comcast
Robert F. Cruickshank III, Ph.D., CableLabs
Derek DiGiacomo, SCTE
Ken Gilbert, P.E. WES.net
Mike Glaser, Cox Communications
Brian T. Patterson, EMerge Alliance
John W. Teague, WES.net
Mark Welsko, P.E., WES.net

Table of Contents

Title	Page Number
1. Introduction	4
2. 10G and the Aging Power Grid	5
2.1. 10G Defined	5
2.2. Aging Infrastructure and the Grid	5
2.3. Reliability & Resiliency (eg: Ca. Public Service Power Shutoff (PSPS Program!))	6
3. Microgrid Defined and the role of alternative energy resources	9
3.1. What is a microgrid?	9
3.2. Simplified Definition	9
3.3. Distributed Energy Resources (DER)	9
3.4. Alternative versus Renewable Energy Resources.	10
3.4.1. Alternative Energy	10
3.4.2. Renewable Energy	11
3.4.3. Clean Energy	11
3.5. Bringing Alternative Energy Resources Together = Microgrid.	12
3.6. What's In It for the Cable Industry/Cable Operators?	15
4. MicroGrid use cases and drivers for the Cable Industry	16
4.1. Edge Facilities and Outside Plant	16
4.2. Electric Vehicles (EVs)	17
4.3. Monetizing Optimum Load Shaping and other EV-Based Grid Support Services	17
5. The Developing Microgrid Technologies and Industry Trends	18
5.1. New term: "Transactive Energy"	18
5.2. Transactive Energy's Potential Benefits to Consumers	19
5.3. The Natural Progression Towards Independent Energy Sourcing and Control	20
6. Conclusion	20
Abbreviations	22
Bibliography & References	23

List of Figures

Title	Page Number
Figure 1 – Simplified Electrical Grid	5
Figure 2 – Growing Risks to the Electrical Infrastructure	7
Figure 3 – World Energy Consumption	8
Figure 4 – North American Power Outages	8
Figure 5 – Typical System Configuration	12
Figure 6 – Microgrid Topology (1)	14
Figure 7 – Microgrid Topology (2)	14
Figure 8 – Microgrid Topology (3)	15
Figure 9 – Critical Infrastructure Facilities	16
Figure 10 – Modernized Electrical Grid	19
Figure 11 – Steps to Transactive Energy Market and Critical Infrastructure Resiliency	20
Figure 12 – Summary: Microgrid Use Cases, Scenarios & Value for the Cable Industry	21

List of Tables

Title	Page Number
Table 1 – Infrastructure Operational Functions	12

1. Introduction

This paper is a joint effort by members of the SCTE Alternative Energy / Microgrid Working Group. The SCTE Alternative Energy / Microgrid Standards Working Group's (AE/MGWG) charter is to "educate and inform the SCTE community on the applicability and use of alternative energy & microgrid technology in cable operator facilities." This includes: defining operational practices and standards; demonstrating the technology is deployable and manageable for service providers; facilitate communication between service providers, industry partners and other standards organizations; and creating a library of microgrid use cases showing how resiliency can be improved, operational costs reduced, and deployment times decreased through the appropriate application of these technologies.

A microgrid is an electrical system that connects multiple sources and loads that is controllable by the user to allow independent operational choices. Currently, some basic alternative energy and microgrid technology have been deployed throughout the cable industry. However, the industry has not yet taken full advantage of existing, available, and relevant advanced powering technologies. Most existing power systems are not ready to work like advanced microgrids. That said, cable operators can, and should, continue to leverage already deployed technologies and test the new approaches to powering. This paper will attempt to address how microgrid technology has continued to evolve, along with the issues facing the application of future microgrid technologies, to illustrate the benefits of adopting a proactive rather than reactive microgrid implementation strategy.

Today, traditional deployments of energy infrastructure in the cable industry includes Direct Current (DC) power plants and Alternating Current (AC) uninterruptable power systems, long term battery storage, transfer switches and switch gear. It also includes generator sets, renewables and other power sources that have been combined in a traditional manner to provide resiliency and sustainability when grid power is lost. While these are many of the basic elements of a microgrid, they often lack the topology and controls required for full microgrid implementation and performance. However, existing deployments are capable of providing a foundation for transition into a more resilient and functional microgrid architecture.

The fundamental premise behind the deployment of a true microgrid architecture by a cable operator is the increasing opportunity to diversify sources of power and therefore enable a more resilient service offering to customers. This would also allow new capital models and power system topology designs to reduce cost of ownership.

2. 10G and the Aging Power Grid

2.1. 10G Defined

The cable industry is quickly moving the needle of network service offerings. The collective initiative was announced as 10G at the 2019 CES show in Las Vegas. This collective push to offer a symmetrical 10 gigabit network to subscribers will enable new ideas, businesses, and things not yet imagined coming to life. “The 10G platform is a combination of technologies that will deliver internet speeds 10 times faster than today’s networks and 100 times faster than what most consumers currently experience. Not only does 10G provide faster symmetrical speeds, but also lowers latencies, enhanced reliability and better security in a scalable manner.”^[1] Powering will continue to be an important component to this developing program, and just like the network will evolve, power strategies should evolve as well.

2.2. Aging Infrastructure and the Grid

The U.S. electric grid (“the grid”) constitutes a vital component of the nation’s critical infrastructure and serves as an essential foundation for the American way of life.

America’s economy, national security and even the health and safety of our citizens depend on the reliable delivery of electricity. The U.S. electric grid is an engineering marvel, with more than 9,200 electric generating units having more than 1 million megawatts of generating capacity connected to more than 600,000 miles of transmission lines, according to the U.S. Department of Energy, Office of Electricity.

This “grid” feeds consumers through an intricate network of transmission lines, substations, distribution lines, and transformers, as shown in Figure 1.

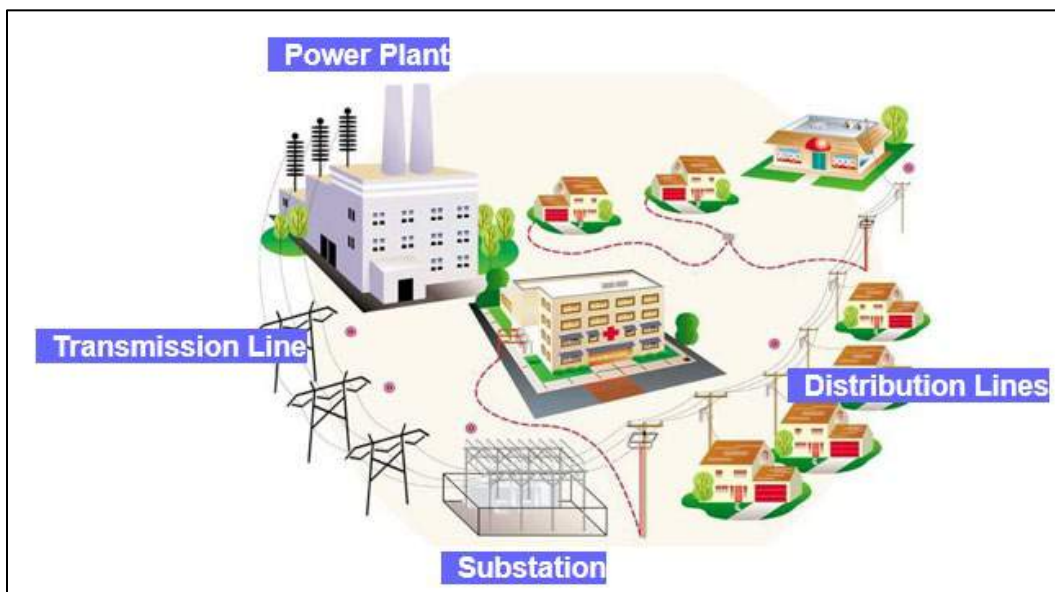


Figure 1 – Simplified Electrical Grid

It is generally understood in the industry that the electrical grid is aging. A good assessment is provided by the American Society of Civil Engineer’s (ASCE) 2017 Infrastructure Report Card, which described most of the U.S.

energy system as predating the turn of the 20th century, with most transmission and distribution lines at full capacity. These systems were constructed in the 1950s and 1960s with a 50-year life expectancy. The annual number of power outages continue to increase, and the ASCE gave the system a D+, the same rating given to it in the previous three 4-year report card cycles, indicating no changes other than the system continues to age.

As a significant portion of the grid was developed and built half a century ago, it did not include the present-day demand for:

- Higher quality power
- Integration of clean, variable renewable, electric vehicles, and distributed energy technologies
- Remote control and data gathering
- Enabling consumer participation
- Higher security and protection from vandalism, terrorism, and weather

The grid was built to simply deliver power with ease of operation, economically, efficiently, and reliable for the age.

This grid connects numerous utilities which are individually focusing on replacement and upgrade as priority over other issues such as aging workforce, regulatory models, and stagnant growth. Meanwhile, they are developing and improving processes and models to manage their assets more efficiently such as:

- New and innovative testing methods, which help to identify and prioritize old equipment that is most in need of repair and/or replacement
- Cable injection and treatment programs
- Breaker refurbishment and upgrade programs
- Wood pole and tower structure testing, treatment, and replacement

Unfortunately, these programs allow the basic infrastructure to continue to age.
















2.3. Reliability & Resiliency (eg: Ca. Public Service Power Shutoff (PSPS Program!))

Cable operators are observing compound annual growth rates (CAGR) of 40 – 50% downstream and 20 – 30% upstream driven by streaming video (including 4K content); newer, delay-sensitive gaming applications; and a general increase in consumption. Over the next few years, cable operators will be faced with numerous decisions in order to meet the exponentially growing needs of their customers.

The upcoming implementations of 10G networks, which promise 10X the bandwidth, assure that there will be an increased need for reliable energy to drive these more powerful devices as they are deployed across the network. This increased demand for powering cable plants and networks comes at a time when traditional utilities are grappling with a myriad of issues, including the fact that they have failed to keep up with the increasing need for truly resilient energy across their mostly centralized networks. This can pose another serious challenge that cable operators will need to contend with.

What's more, to contend with existing and up-and-coming competitors, cable operators are also increasing their dependence on reliable and consistent power to service newer offerings that are being driven by the proliferation of Internet of things (IoT) applications, including in-home security systems and other "smart devices." This all increases backhaul requirements and intensifies the demand for resilient and reliable power.

These dynamics are forcing operators to reconsider their investment strategies and the soundness of making longer term capital expenditures to reduce risk, reduce costs and remain competitive. Microgrids may provide a viable strategy, by reducing the dependency on traditional power providers with the possible concurrent benefit of reducing energy-related operating expenses over time. Fortunately, many cable operators have been quietly engaged in building out what could be considered microgrids of sorts. By installing renewables and energy storage systems, operators are adding resiliency and maintaining operations in the event of power failures resulting from disruptions in the power grid. An additional benefit of these hybrid microgrids can also be realized by reducing demand charges and possibly reselling power to utilities under power purchase agreements. Other drivers of the trend to increase independence from traditional providers are the increased intensity, regularity, and length of weather disruptions, and the inability of utilities to easily make enhancements to their generation and distribution capabilities. Issues related to localities objecting to pipeline placement, new power line installations and other regulatory challenges are adding to risks associated with dependency on traditional utility operators. Operators in California have recognized the potential for such grid disruptions and have begun to respond by promoting centers in areas less likely to be impacted. Figure 2 provides a list of electrical infrastructure hazards and associated risks.

Natural Hazards	Human Hazards	Operational Hazards
 Ice, snow and extreme cold weather	 Physical attacks	 Geomagnetic and electromagnetic pulses
 Thunderstorms, tornados and hurricane-force winds	 Cyber attacks	 Aging infrastructure
 Storm surge, flooding and increased precipitation	 Workforce turnover and loss of institutional knowledge	 Capacity Constraints
 Increasing temperature and extreme hot weather	 Human Error	 Dependencies and supply chain interruptions
 Earthquakes		 Inherent instability from renewable resources
 Wildfires		

Adapted from Argonne National Laboratory, 2016

Figure 2 – Growing Risks to the Electrical Infrastructure

Growing energy demand is also impacting the existing industry growing risk situation. Despite increased efficiency, the US Energy Administration projects that world energy demand will increase by 28% by 2040. Much of this demand is driven by the increasing digitization of society, data centers, telecommunications systems,

and our households, that continue to require more energy to operate the devices we have become so dependent upon.




Figure 3 – World Energy Consumption

Also compounding the problem is the increased intensity and frequency of major storms, fires and other natural disasters. Outages of more than 24 hours impact every citizen. Food supply, manufacturers, communications networks, and simple changes to every day routine have consequences beyond financial. Also to be considered are the financial and manpower resources required to respond to such events, the pressure on traditional utilities to modernize generation and distribution networks, and the impact on our communities.




Figure 4 – North American Power Outages

The demand for quality, dependable power is also extending beyond primary cable plant and includes the growing demands of peripheral (edge) requirements for power as the technology evolves. Utility costs for these edge sites is also increasing along with the risks of losing utility power. Power fundamentals such as native DC vs. AC supply of energy also adds to the cost and complexity of supporting cable operations. With recent advances in direct current technologies and deployment capabilities DC powered networks can increase system efficiency, particularly as the industry integrates native dc powered renewables and energy storage devices into their power networks.

3. Microgrid Defined and the role of alternative energy resources

3.1. What is a microgrid?

The SCTE Alternative Energy/Microgrid Standards Working Group (AE/MGWG) developed a definition for microgrids as: “Microgrids may be defined as a localized group of interconnected and managed electricity sources, storage and loads that can be connected with other local microgrids and/or the traditional electrical utility grid (macro grid) but can seamlessly and selectively disconnect from them and function independently as conditions, policies or economics dictate.”

3.2. Simplified Definition

A more simplified definition of a microgrid can be stated as: A microgrid is an electrical system that connects multiple sources and loads that is controllable by the user to allow independent operational choices.

3.3. Distributed Energy Resources (DER)

The United States electric grid is comprised of bulk power generation, distribution, transmission, and consuming entities. The infrastructure enabling just about everything in our modern daily lives has remained largely consistent for close to one hundred years. AC power is generated, moved in higher voltage from generation to substations, and voltage is reduced to feed transmission and then ultimately converted into the low voltage AC power typically used at commercial and residential facilities.

With the advent of solar and what we refer to as “alternate energy” sources, the model of bulk power sources begins to be challenged. The traditional grid model is one of stability, predictability, and regulation to uphold high expectations of availability. To satisfy these expectations of high availability the current system is made up of large central base-load fossil fueled power generation plants and supplemented by addition fossil fueled based peaker plants when demand increases beyond base load capabilities or are needed to supplement when base load system is unavailable. When we plug something into an electrical outlet here in north America, there is no second-guessing power availability. Distributed energy resources can become a viable model to anyone wishing to pursue alternatives to the traditional power grid.

The new challenge of incorporating solar photovoltaic (PV) systems, wind resources, storage (like batteries) raises electrical engineering questions as well as financial market questions. DERs can be defined as: “Distributed energy resources are small, modular, energy generation and storage technologies that provide electric capacity or energy where you need it. DER systems may be either connected to the local electric power grid or isolated from the grid in stand-alone applications. DER technologies include wind turbines, solar/photovoltaics (PV), fuel cells, microturbines, reciprocating engines, combustion turbines, cogeneration, and energy storage systems.”^[2] Cable operators are in a very good position to evaluate and deploy DERs due to hybrid fiber coax (HFC) communication

network architectures. Critical facilities, people space and even outside plant can be considered for distributed energy adaptation.

There are several use cases that can promote the deployment of a DER for cable operators. These primary cases include lowering grid dependency, becoming less utility grid dependent, and enhancing availability of power. In locations where peak demand charges become very high during seasonal changes (high air-conditioning or heating needs), a local DER such as a grid-tied PV plant can help reduce excess charges. The PV plant can be engineered to match the forecasted spikes in utility grid demands that associate with a higher billing to the cable operator. In 2019, California power providers began to institute public safety power shutdown (PSPS) events. Weather conditions supporting high fire risks have required grid providers to turn off supply in order to lower risks of fire. Communication providers need to become less dependent on the grid to ensure service availability. DERs can help address that need. This approach is naturally related to the third use case of enhancing power availability. NREL conducted a backup power study for cable operators and that report can be found in Reference [3].

Distributed energy resources can play an important role in cable's infrastructure. Determining the strategic adoption can be done following a few steps. First, identify the electrical requirement and/or problem. Are the bills too high? Is power availability becoming unacceptable? Are there renewable energy requirements needed for positive marketing? Does run-time in the absence of grid or traditional generator backup need to be extended? After the applicable question/s is/are identified and answered, the technology can be identified to address the problem. PV has been a reliable go-to; however, each situation will dictate what power generation source would fit best. Second, approach the local utility provider to discuss incentives. If the local power provider has a mandate to lower loads or incorporate renewable power sources, there may be financial awards to help offset DER deployment costs. Third, build out the project plan to commissioning. As with any change to day-to-day operations, a DER implementation requires a solid project plan to ensure successful incorporation into existing network topologies. Finally, document post commission needs, such as maintenance, system milestones and support providers. Cable operator infrastructure requires working with a multi-vendor ecosystem and the DER turn-up can introduce new providers of power service if the existing pool of resources do not specialize in essential power source management. This relationship will be especially important as a DER investment can extend beyond 20 years.

3.4. Alternative versus Renewable Energy Resources.

The differences between “alternative energy,” “renewable energy,” and “clean energy,” might not be obvious. Each term is unique and has its own individual definition.

3.4.1. Alternative Energy

Alternative energy refers to sources of usable energy that can replace conventional energy sources (usually, without undesirable side effects). The term “alternative energy” is typically used to refer to sources of energy other than nuclear energy or fossil fuels.

Throughout the course of history, “alternative energy” has referred to different things. There was a time when nuclear energy was considered an alternative to conventional energy and was thus called “alternative energy.” The term is ever evolving.

Today, a form of “alternative energy” might also be renewable energy, or clean energy, or both. The terms are often interchangeable, but not the same.

3.4.2. Renewable Energy

Renewable energy is any type of energy which comes from renewable natural resources, such as wind, rain, sunlight, geothermal heat, and tides. It is referred to as “renewable” because it does not run out.

People have begun to turn to this type of energy due to the rising oil prices, and the prospect that one day sources of fossil fuels may be depleted. Also, concerns about the adverse effects that our conventional energy sources have on the environment have played a big role in the advancement and adoption of renewable energy sources.

Among the different types of renewable energy, wind power is one which is growing in its use. The number of users who have some form of wind power installed has increased, with the current worldwide capacity being about 100 GW. In addition, the traditional power grid has leveraged commercial wind farms to diversify power sources in environmental correct locations such as the mid-west and parts of Texas.

3.4.3. Clean Energy

“Clean energy” is simply any form of energy which is created with clean, harmless, and non-polluting methods. Most renewable energy sources are also clean energy sources, but not all.

One such example is geothermal power. It may be a renewable energy source, but some geothermal energy processes can be harmful to the environment. Therefore, this is not always a clean energy. However, there are also other forms of geothermal energy which are harmless and clean.

Clean energy makes less impact on the environment than our current conventional energy sources do. It creates an insignificant amount of carbon dioxide, and its use can reduce the speed of global warming – or global pollution.^[4]

3.5. Bringing Alternative Energy Resources Together = Microgrid.

Typical cable industry critical infrastructures (CIs) are constructed with power systems that contain emergency back-up systems, such as fossil fueled generators and lead acid battery-based DC power plants, to ensure system reliability during utility power outage events. A typical system configuration example is shown below in Figure 5.

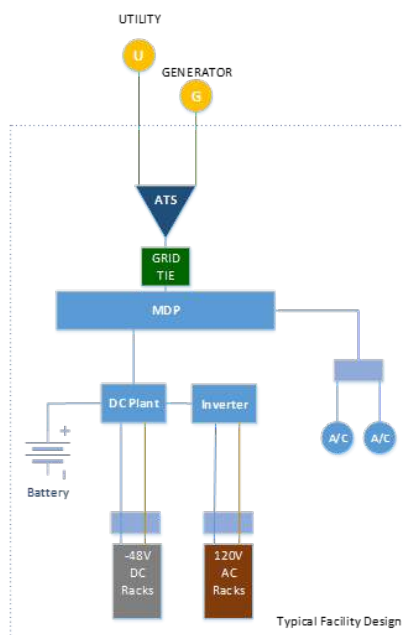


Figure 5 – Typical System Configuration

The current state of this infrastructure is characterized with specific technologies characterized by the functions and related adjacent operational functions as described in Table 1.

Table 1 – Infrastructure Operational Functions

Business Function	Technology	Technology Function	Operational Considerations
Primary/Secondary Power Source	Utility Service and Generator	To power the critical load	<ul style="list-style-type: none"> • Site is dependent upon reliable utility service • Site operates on either utility OR generator • A brief outage is required for generator start-up • Most generators are fossil fuel based • Many sites with dual generators
Source Transfer	Automatic Transfer Switch (ATS)	Monitor utility power and controls	<ul style="list-style-type: none"> • No parallel operation • No closed transition transfers

		transfer of power source	
Power Distribution	<ul style="list-style-type: none"> • -48v DC/battery plant • Dual bulk feeds to critical loads 	Service continuity during transition from utility to generator	<ul style="list-style-type: none"> • High initial investment with stranded capacity costs • Assorted power distribution configurations • System changes and adds are costly with specialized labor and material which can influence time • Power density varies widely but initial build investment needs to anticipate worse case loads • Inverters can allow for elimination of UPS
	Inverter	Critical AC loads through inverter (eliminating UPSs in process)	
Stable Operating Environment	HVAC		<ul style="list-style-type: none"> • Backed up by generator for power during commercial power outages

Bringing a variety of power sources together to feed and support a load or multiple loads is a microgrid. The cable industry, to some extent, has a head start in the progression path to microgrids that should be evident from the above. Some examples of microgrid topologies are shown on the following page.

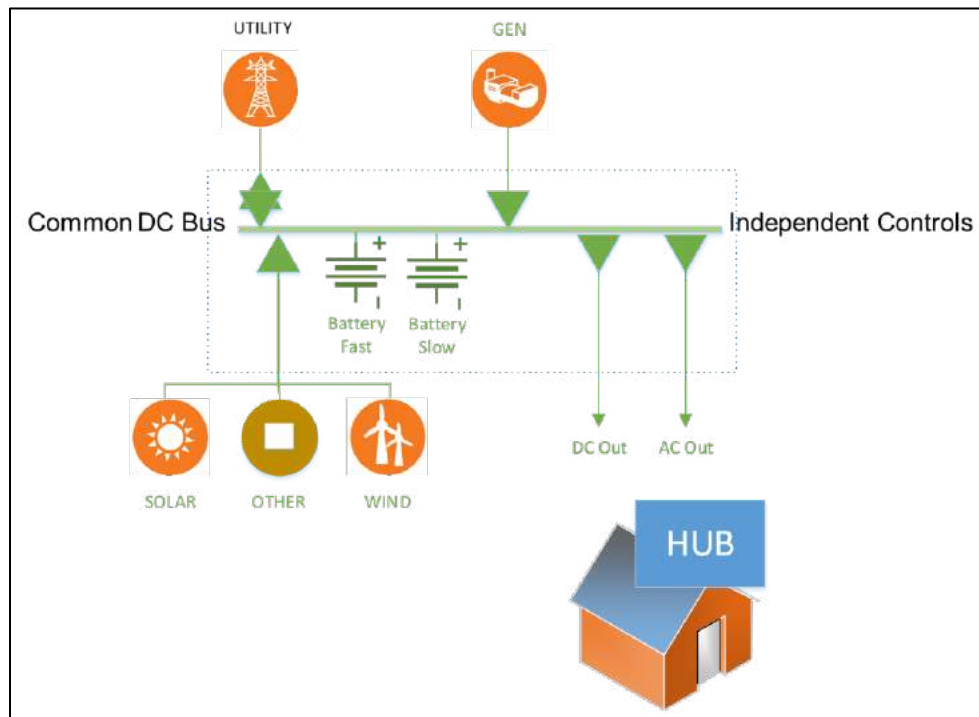


Figure 6 – Microgrid Topology (1)

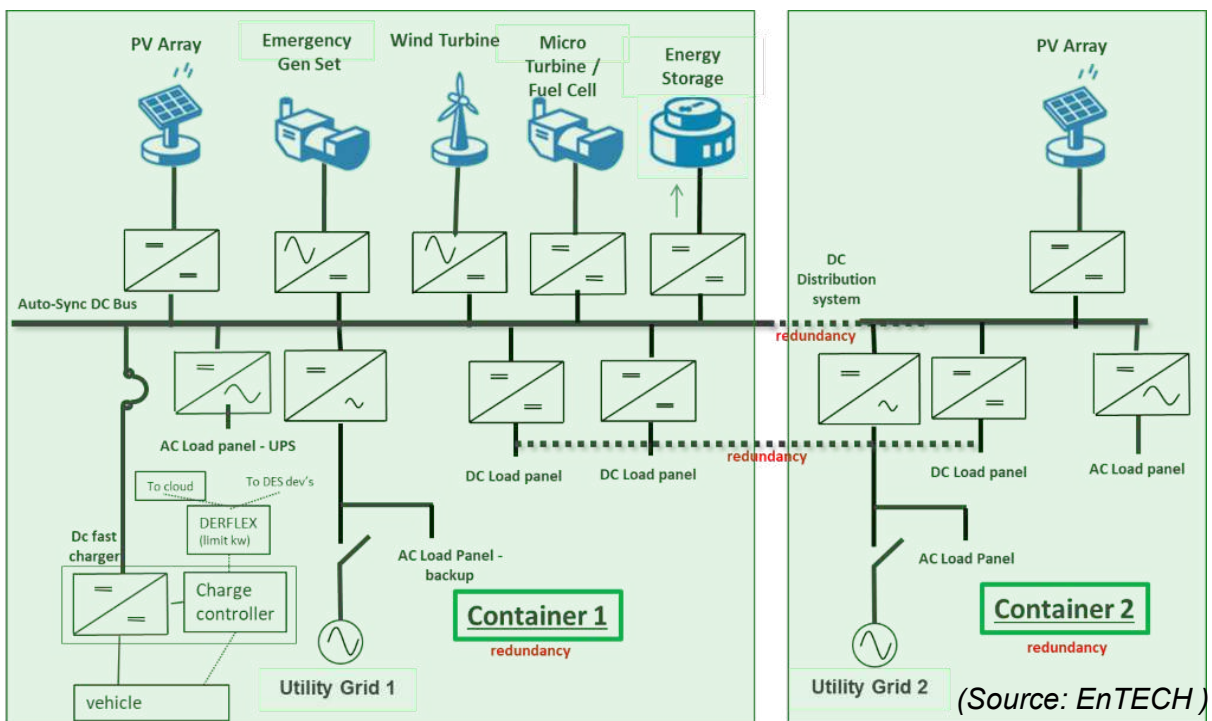


Figure 7 – Microgrid Topology (2)

The preceding configurations use a DC power collector bus topology. An example of an AC power configured microgrid is shown in Figure 8 (Source: Schneider Electric).

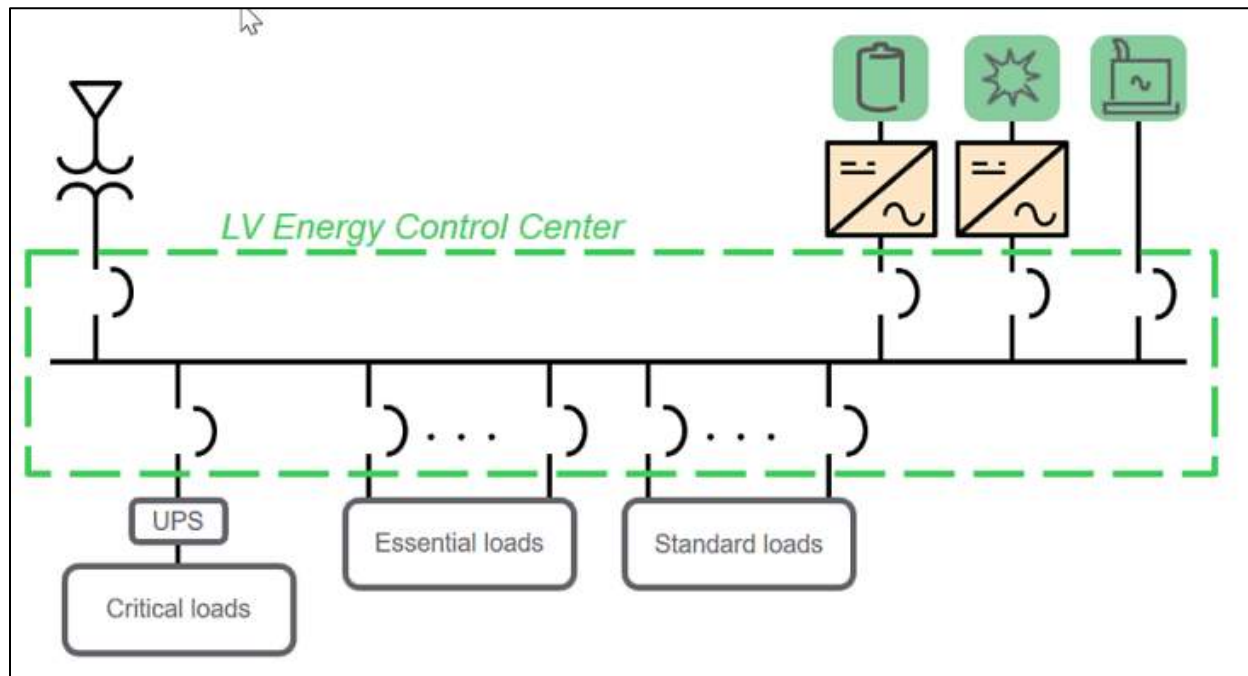


Figure 8 – Microgrid Topology (3)

3.6. What's In It for the Cable Industry/Cable Operators?

The DOE and other government agencies have been spearheading the definition of microgrids while documenting their value. Now, with the active involvement of two predominant industry organizations - NREL (National Renewable Energy Laboratory) and CableLabs - this should send a clear message that it is time for the cable industry to grab the lead, creating microgrid standards specifically designed for cable operators.

As introduced earlier in the paper, microgrids are very important to the cable industry for several reasons:

- 1) They provide additional resiliency and reliability in times of unplanned power outages. During severe weather, microgrids can provide the ability to continue providing power to critical facilities.
- 2) With competitively priced self-generation from renewable and other energy sources, microgrids can provide a hedge to increasing costs of energy for cable operations.
- 3) Coincident with time-of-use electricity pricing, microgrids can allow cable operators to buy low and sell high for their operational energy needs, as well as the energy needs of their customers.
- 4) Microgrids provide operational independence from local utilities allowing control and usage of dispatchable power sources independent of local utility performance.

Operational cost containment is an important practice for cable operators, and power is no exception. The idea of purchasing energy at low rates and selling high over a 24-hour period will become increasingly important. Time-of-use rates become pervasive to effectively manage the supply and demand balance of the grid in the presence of increasing renewable energy resources. This is evidenced by the increasing number of cities, states and countries

that have made 100% renewable generation commitments, most notably, the states of California, Hawaii and several others by 2045.

While these economic drivers are valid, other areas of the country with lower electrical costs are justifying microgrids by combining other value propositions with Time-of-use gains. Such benefits include added resiliency, demand shaving, and the utilization and control benefits of higher voltage direct current in a DC-Coupled microgrid.

Basic microgrid technologies are currently deployed within the cable industry, but they are not positioned to best leverage traditional microgrid techniques. The opportunity exists for the industry to take stock of what is changing in their networks, followed by an exploration of how the industry should or can leverage new technologies (including power). As the industry strives to enhance the customer experience, there needs to be a conscious review of how energy use has evolved. As new cable technologies are deployed in the access networks, infrastructure costs to accommodate these changes could end up being fiscally prohibitive in meeting business goals without proper consideration of modular microgrid technology deployments.

The regionalized nature of the cable industry is conducive to microgrid developments. The cable industry could develop microgrids and sell power to/within itself as well as sell power to outside partners. The steadily improving return on investment (ROI) related to microgrid and alternative energy technologies, coupled with regionalized opportunities that exist for the cable industry/cable operators, make a valid argument that there are tangible benefits for microgrid development.

4. MicroGrid use cases and drivers for the Cable Industry

4.1. Edge Facilities and Outside Plant

It is well documented by SCTE, and as depicted in Figure 9, that the majority of the cable industry’s critical infrastructure power footprint is housed within edge facilities and outside plant.

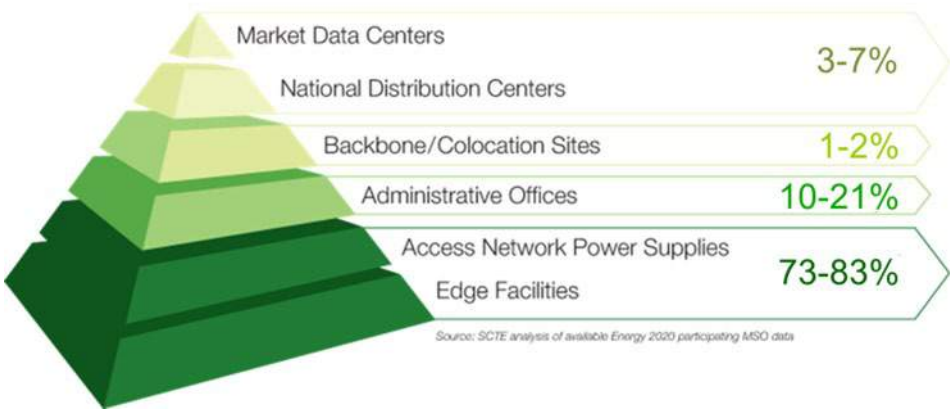


Figure 9 – Critical Infrastructure Facilities

Outside Plant Infrastructure: Unique to cable and telecom providers is the outside plant and especially the hundreds of thousands of power supplies enabling the hybrid fiber coax network (HFC). Typically, power is provided to the center conductor of the HFC plant via a local power supply. The power supply receives power

from the utility provider typically on the very same pole the cable infrastructure resides. What kind of creative approach could be taken to address the diverse nature of outside plant powering? Self-generation of power in strategic places could be an opportunity for investigation. New innovative power supplies and infrastructure are being investigated and new standards are in process to allow this innovation. New energy storage technologies will also play an important role.

Edge Facility Infrastructure: According to ANSI SCTE 226 2015, Class B-D facilities are smaller in nature, support, and design criteria than the larger Class A data center centric facilities. This does not mean that the importance of these infrastructure rich points of presence are small. Evaluation of microgrid deployment options are not easy to determine. Careful examination of each facility should be conducted, and a reusable model built to evaluate the benefit and application of microgrid options at the smaller facilities. The model needs to consider multiple data points such as downtime, cost of power, number of subscribers, and types of subscribers (commercial vs. residential). Network architecture dependencies are important factors when determining microgrid investment. All new infrastructure power solutions must be modular and scalable.

4.2. Electric Vehicles (EVs)

The cable industry maintains a fleet of nearly a quarter million vehicles worldwide. In recent years, the industry often encourages technicians to overnight company vehicles at their homes. Continuing to do so will create unique financial opportunities to leverage the industry's broadband infrastructure to manage the charging of cable's future fleet of EVs—and perhaps even private vehicles—across small and large geographic areas.^[5]

EVs and hybrid EVs will play an increasingly important role in cable operations. Due to technological advances and continuing declines in battery costs, personal EVs provide fuel savings of nearly \$1,000 per year and EVs overall are cheaper than equivalent combustion-engine models for many applications. Deployment of EVs are increasing every year and should be part of microgrid applications and growth. This will help fill the need for increased power reliability and resiliency.

Global decarbonization is driving the further electrification of our world and this will result in increased electricity consumption of 38% by 2050^[6]. Electricity is rising in popularity as it is easier to transport, deliver, store, and use. Society's growing dependence on electric power is resulting in an exponential rise in consumption in applications such as data centers and EVs. In and of themselves, EVs are likely to create 20%-30% additional load on the electric grid^[6], helping to make the case for the optimization of charging strategies. As such, Time-of-use pricing of electricity is legislated in several states and, along with variable pricing is expected to be pervasive.^[7] The departure from fixed electricity rates raises crucial financial questions of 1) when to charge EVs based on varying electricity cost, and 2) how to enable the cable industry to specify methods that monetize the value stack created by managed fleet vehicle charging. As mentioned earlier, cost management of operations for cable operators is important to the business.

4.3. Monetizing Optimum Load Shaping and other EV-Based Grid Support Services

Several technologies and market trends are at play that can both negatively and positively affect the cable industry in terms of energy cost. With the increased deployment of EVs, the production, distribution, and use of electricity is rapidly evolving for the charging infrastructure that will be needed, creating critical functionality gaps in managing the grid. EV charging is already stressing grid capacity and affecting the cost of electric power. Yet, when networked and managed in a coordinated fashion, batteries and EVs are proving their ability to provide grid support services such as load shaping, peak reduction, and active power quality management via reactive power and frequency support. Together, grid support services from EVs can create a value stack consisting of reductions in operational costs, maintenance, and new construction of power plants, transmission, and distribution facilities.

If cable operators can reap the rewards of selling power via a microgrid and leverage the collective energy storage across a substantially large EV fleet, the benefits could be substantial. Communications and controls will play an important role in such a scenario.

Since 1882, the grid has operated such that supply from power generators anticipates and follows the demand for electricity. In recent years, the continuing decline in the levelized cost of energy from wind and solar power is such that, in much of the world, construction and operation of new renewable energy sources are less costly than the ongoing operation of existing fossil-fueled power plants. Yet, the demand for electricity is often not coincident in time with the supply from renewables, which themselves are variable and not dispatchable (i.e., not controllable). As such, widespread, pervasive coordination of demand will be valuable in orchestrating electric loads, such as EV charging, to follow the least costly forms of fossil-based and renewable supply.

A growing body of research from several U.S. Department of Energy National Laboratories and throughout industry indicates that load shaping will be increasingly important in reducing the cost of operations in the grid, microgrids, and nanogrids.^[8] What is missing are methods to optimally shape load based on the holistic consideration of generation, distribution, and storage. Optimum Load Shaping (OLS) is a newly developed software-based proprietary technology to minimize power generation costs and carbon dioxide emissions that informs electrically powered devices of the forecast times of the lowest cost and cleanest supply. OLS uses an end-to-end generation-to-load algorithm that jointly optimizes supply (part 1) and demand (part 2). Efforts are underway in the SCTE Energy Management Subcommittee, Microgrid Working Group, to propose a standard that will specify the creation of an OLS, its transmission across different networks, and the actions taken by a receiving device, such as an EV or battery charger, that modifies its behavior to minimize differences over time from the OLS.

Forecast Optimum Load Shapes (OLs) can help monetize the cable industry's future fleet of electric vehicles and facility batteries to provide the critically needed end-to-end, generation-to-load control of the electric power grid. An OLS provides grid control and consists of a set of numbers (e.g. target load for hours 1-24) that forecasts the most efficient electrical supply in grids, microgrids, and nanogrids, so that all stakeholders - generation entities, utilities, distributors, retailers, and consumers—can reduce their electricity costs and carbon emissions.

5. The Developing Microgrid Technologies and Industry Trends


As microgrids are developed in greater and greater numbers, microgrid interconnections will also increase. Interconnections will occur directly between adjacent microgrids, but in most cases will be interconnected through the existing electric utility infrastructure. Microgrids and the interconnection of microgrids require new ways to track energy use, power generation by distributed energy resources, and load control in a way that transactions between microgrid owners and operators are being developed. These situations and ongoing engineering developments are what make up the new “Transactive Energy” evolutions as described below.

5.1. New term: “Transactive Energy”

The increased use of renewable energy and distributed energy management technologies offers the potential for significant efficiency improvements through market-based transactive exchanges between energy producers and energy consumers.

To enable these sorts of exchanges, however, the modernized grid will require new economic tools and processes. “Transactive energy” is the broad term used to describe this new approach and can be defined as *“a system of*

economic and control mechanisms that allows the dynamic balance of supply and demand across the entire electrical infrastructure using value as a key operational parameter.”



Source: EPRI

Figure 10 – Modernized Electrical Grid

5.2. Transactive Energy’s Potential Benefits to Consumers

The following list outlines the benefits of moving to a more transactive-based energy model:

- Better utilization of grid assets: Everything from transformers and switches, to vehicle-charging stations and smart meters, can lower costs when optimized, especially during peak demand conditions.
- Greater resilience and reliability: During large storms, a reduced length and frequency of outages
- Greater control over personal energy use: Empowerment of choice and information provided to consumers
- Increased use of renewable energy resources: Gives individual consumers the satisfaction of contributing to larger societal environmental sustainability goals

The growth of a cable operator’s business intelligence that is required for this evolution, including the development of artificial intelligence (AI), will evolve into a big data management platform. This platform could make for a new business opportunity for cable operators as well as provide market participation opportunities in various regions of the country.

5.3. The Natural Progression Towards Independent Energy Sourcing and Control

Originally driven by economics (cost savings) and social responsibility (energy management), the first step is the installation and application of renewable energy technologies and systems. The next steps are the progression paths to the future with a corporate version of self-actualization (maximizing potential) by participating in the energy market as both a supplier and consumer, with the ability to self-determine the most appropriate energy source based upon real-time availability and needs. The figure below describes the steps in the progression path to the Transactive Energy future

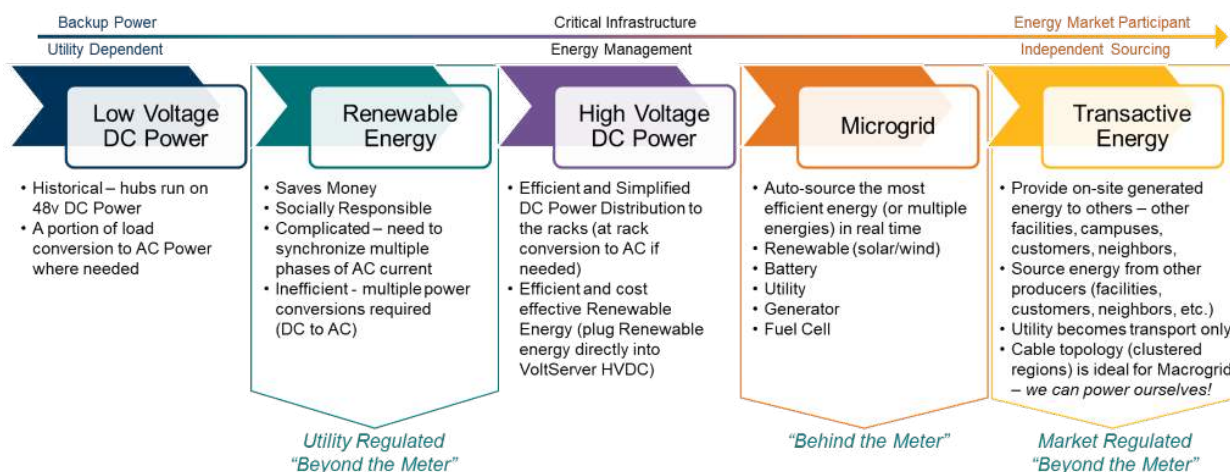


Figure 11 – Steps to Transactive Energy Market and Critical Infrastructure Resiliency

6. Conclusion

As the cable industry looks to further standardize on strategic energy concepts like microgrids, opportunity awaits for development, trial and ultimate deployment of microgrid technologies. It has been shown that microgrids are being deployed within many industries and use cases globally. Economics, the rising cost of electricity and the focus on a more sustainable future are the drives for the widescale deployment of solar, wind and other distributed energy resources. New technology large scale energy storage is also becoming more cost effective and is playing a bigger and bigger roll in the new energy future. This paper characterized why it's a sound ideae to fully reap the benefits of multi-source smart power grid scenarios, adhering to cable specific use cases mentioned in the early sections of this paper. Finally, the consensus body represented in the active SCTE working group will pave the way to help accelerate the deployment of robust microgrid solutions. These will enable a foundation of new network ideas, as envisioned and embodied in the announcement of 10G, early in 2019. Microgrids will play a big roll in powering 10G with a summary of use cases, use scenarios and values to the industry in figure below.

Summary: Microgrid Use Cases, Scenarios and Values



Use Cases

1. Market Data Center
2. National Distribution Center
3. Backbone / Colocation Sites
4. Administrative Offices
5. Access Network Power Supplies
6. Edge Facilities
7. Outside Plant
8. Fleet – EV Charging



Value Propositions:

What problem are we trying to solve?

1. Energy Efficiency
2. Resiliency
3. Reduce utility grid Dependency
4. Sustainability goals
5. Auto-source of lowest cost generation

Microgrid Use Scenarios

1. Demand Response
2. Peak Shaving
3. Black Start
4. DER flexibility
5. Islanding
6. Parallel
7. Transactive Energy
8. Improved reliability w/multiple sources of electricity
9. Many more

Figure 12 – Summary: Microgrid Use Cases, Scenarios & Value for the Cable Industry

Abbreviations

10G	10 gigabits per second
AC	alternating current
AE/MWG	Alternative Energy / Microgrid Working Groups
AI	artificial intelligence
ANSI	American National Standards Institute
ASCE	American Society of Civil Engineers
ATS	automatic transfer switch
CAGR	compound annual growth rate
CES	Consumer Electronics Show
CI	critical infrastructure
CO ₂	carbon dioxide
DC	direct current
DER	Distributed Energy Resource
DOCSIS	Data Over Cable Service Interface Specification
DOE	Department of Energy
EV	electric vehicle
FEMP	Federal Energy Management Program
GMLC	Grid Modernization Laboratory Consortium
GW	gigawatt
HFC	hybrid fiber coax
HVAC	heating, ventilation and air conditioning
IoT	Internet of Things
ISBE	International Society of Broadband Experts
NREL	National Renewable Energy Laboratory
OLS	optimum load shaping
PSPS	public safety power shutdown
PV	photovoltaic
ROI	return on investment
SCTE	Society of Cable Telecommunications Engineers
UPS	uninterruptible power supply

Bibliography & References

- [1] *Cable Labs*. <https://vimeo.com/309550130/21b19612e4>
 - [2] “*Using Distributed Energy Resources*”, *Federal Energy Management Program (FEMP)*, 2013. <https://www.nrel.gov/docs/fy02osti/31570.pdf>
 - [3] *SCTE Journal of Energy Management*, V2 N2, 2007. <https://www.nrel.gov/docs/fy17osti/69034.pdf>
 - [4] <https://livingclean.com/differences-alternative-renewable-clean-energy>
 - [5] <https://sepapower.org/resource/a-comprehensive-guide-to-electric-vehicle-managed-charging>
 - [6] “*Electrification Futures Study*”, National Renewable Energy Lab, TP-6A20-71500, 2018
 - [7] “*Residential Electric Vehicle Time-Varying Rates That Work: Attributes That Increase Enrollment*”, Smart Electric Power Alliance, November 2019. <https://sepapower.org/resource/residential-electric-vehicle-time-varying-rates-that-work-attributes-that-increase-enrollment>
 - [8] “*Survey of Distributed Energy Resource Interconnection and Interoperability Standards*”, U.S. Department of Energy, Grid Modernization Laboratory Consortium (GMLC), 2020
- “*Economic Benefits of Increasing Electric Grid Resilience to Weather Outages*”. US Council of Economic Advisers and the U.S. Department of Energy’s Office of Electricity Delivery and Energy Reliability, August 2013
- “*America’s Infrastructure Grade*”. www.infrastructurereportcard.org. American Society of Civil Engineer’s, 2017
- SCTE 226 2015 Cable Facility Classification Definitions and Requirements

Profile Management Informed Proactive Network Maintenance

Comparison of RxMER Per Subcarrier, Bit Loading, and Impairment Driven versus Measurement Variability

A Technical Paper prepared for SCTE•ISBE by

Jason Rupe

Principal Architect

CableLabs®

858 Coal Creek Circle, Louisville, CO 80027

303.661.3332

j.rupe@cablelabs.com

Jingjie Zhu

Senior Engineer

CableLabs

858 Coal Creek Circle, Louisville, CO 80027

303.661.3312

j.zhu@cablelabs.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Approaches Studied.....	5
2.1. Prioritize by worst profile group without anomaly detection.....	5
2.2. Prioritize by worst profile modem without anomaly detection.....	6
2.3. Prioritize by profile bitrate impact using anomaly detection.....	7
2.4. Prioritize by bit-loading or RxMER per subcarrier impact using anomaly detection.....	8
2.5. Prioritize by bit-loading or RxMER per subcarrier without anomaly detection.....	9
3. Details, Discussion, and Evaluation.....	9
3.1. A note on using these results in a PNM framework	9
3.2. A note on clustering work by anomalies	10
3.3. What to monitor	11
3.4. Training data for anomaly detection	11
3.5. Bit loading	11
3.6. Measuring performance	12
3.7. Comparison of measurements	12
3.8. Manual simulation of proactive maintenance impacts.....	22
4. Conclusion and Future Work.....	26
Abbreviations.....	27
Bibliography & References	28

List of Figures

Title	Page Number
Figure 1 - Severity versus normal score from RxMER.....	13
Figure 2 - Severity versus normal score from bit load.....	13
Figure 3 - Severity from RxMER value versus severity from bit loading value	14
Figure 4 - Severity from RxMER value versus severity from PMA value	15
Figure 5 - Severity from bit loading versus severity from PMA value.....	15
Figure 6 - Normal score from RxMER value versus normal score from bit loading value.....	16
Figure 7 - Severity for RxMER value versus bit loading for CM that score a positive PMA severity	16
Figure 8 - Severity versus normal score from RxMER for LTE ingress impairments.....	17
Figure 9 - Severity versus normal score from bit load for LTE ingress impairments.....	17
Figure 10 - Severity from RxMER value versus severity from bit loading value for LTE ingress impairments	18
Figure 11 - Severity from RxMER value versus severity from PMA value for LTE ingress impairments...	18
Figure 12 - Severity from bit loading versus severity from PMA value for LTE ingress impairments	19
Figure 13 - Normal score from RxMER value versus normal score from bit loading value for LTE ingress impairments.....	19
Figure 14 - Severity from RxMER value versus severity from bit loading value for wave impairments	20
Figure 15 - Severity from RxMER value versus severity from bit loading value for rolloff impairments.....	20
Figure 16 - Severity from RxMER value versus severity from bit loading value for suckout impairments	21
Figure 17 - Severity from RxMER value versus severity from bit loading value for spike impairments	21

Figure 18 - Severity from RxMER value versus severity from bit loading value for modems reporting multiple impairment types.....	22
Figure 19 - 526 CMs on a common channel, many experiencing impairments that appear related.....	23
Figure 20 - Cumulative bit load improvements over multiple rounds of impairment removal for both methods.....	24
Figure 21 - Cumulative J value (PMA) improvements over multiple rounds of impairment removal for both methods	24

List of Tables

Title	Page Number
Table 1- Results prioritized by bit load (Severity(Bitload)).....	25
Table 2 - Results prioritized by J value (equivalent to Severity(PMA)).....	25

1. Introduction

With the introduction of DOCSIS® 3.1 technology, profiles and profile management became important to operators who wish to get the most out of their network capacity. With the resiliency advantages of DOCSIS 3.1 technology, network impairments impact service through a loss of capacity but not failure of service until these impairments become severe. Therefore, there is an opportunity to use lost capacity as a way to measure the severity of impairments, and to prioritize proactive network maintenance (PNM) work as well. But with profiles being limited in number on many cable modem termination systems (CMTSs), the theoretical maximum capacity can't always be obtained. So, profiles need to be considered to truly measure the impact of proactive maintenance operations.

In this paper, we present some competing methods for prioritizing PNM work in terms of optimal possible profiles, as well as the measures involved in setting these profiles (RxMER per subcarrier, and bit load). Depending on the conditions of the operator's network and PNM tool capabilities, some solutions will be better than others. So, we will help operators decide an approach that works best for them. The solutions we compare will be offered as workers in our Proactive Operations (ProOps) platform as well, so that anyone interested can conveniently try them for themselves.

Identifying PNM opportunities is one challenge, but prioritizing them is another, and selecting the important problems to work on is yet another. All these steps must be done right for PNM to be effective because repair resources are limited.

Profile management is the practice of optimizing bitrates overall for the cable modems (CMs) on a CMTS; Profile Management Application (PMA) is the name of the application that optimizes profiles, say to maximize overall bit rate across all CMs subject to a limited number of profiles to share.

Operators express that they are reluctant to do proactive maintenance in part because they do not know which problems to address, or which problems will impact service in the future. But the way DOCSIS resiliency mechanisms work, bitrate is sacrificed for service reliability in the face of impairments. So, while one impairment may not itself impact service, a future one will because the impacts pile up until the customer notices. All impairments impact bitrate in some way, be that through codewords that need correction, data that has to be re-sent, or subcarriers that have to be avoided, for examples. This suggests that the impact to bitrate may be a good candidate for prioritizing and selecting important proactive maintenance projects to tackle.

But how do you determine the bitrate impact of impairments? That is where a look at PMA can be informative. For downstream PMA, downstream RxMER per subcarrier is collected and used to determine bit loading, which in turn informs bit load, which informs how the profiles are set. Therefore, there are three potential candidates to consider: RxMER per subcarrier, bit loading, or the profiles assigned.

Therefore, we investigate each of these approaches as methods for assigning PNM work. For some of these methods to work, however, anomaly detection is all but necessary. We also look at possible methods that include anomaly detection using a machine learning approach, and methods that avoid use of an anomaly detector (AD) when possible, for cases where AD is not available. Statistical methods can be used to detect impairments as well; for example, operators and vendors who are members of CableLabs can obtain a copy of Spectra which detects anomalies in spectrum data, including methods that were contributed back by Comcast.

After outlining methods for consideration, we test these methods with an available data set, and compare the merits of the approaches. While we don't have access to operator networks to test the methods in deployment, the comparisons we have made are informative enough to narrow the options down and inform which may be tried first in various situations. We encoded some of the better options into the ProOps platform so we can work with operators to test and tune the methods for their use. These solutions are available to CableLabs member operators for free, and provided in some form to vendors as well.

2. Approaches Studied

Next, we introduce the approaches we defined and compared for prioritizing PNM work. Each of these can be updated after information changes, such as the completion of maintenance.

2.1. Prioritize by worst profile group without anomaly detection

The target in this method is the worst performing CMs in the worst profile. If your PMA calculation uses anomaly detection within it, then anomaly detection is used indirectly but its information is not directly used in this case.

1. For each CM, calculate its profile from PMA based on the configured number of profiles in its CMTS (Assigned Profile, as is typically done with PMA), and the best profile it could have been assigned with unlimited profiles possible. (Possible Profile is a profile it would be assigned if it were in a cluster of one). Note this is before consideration of impairments that can be removed.
2. Calculate the bit rates (bit loading) these CMs should achieve with these two profiles (BR_AP and BR_PP).¹
3. Cluster CMs by profile. Given the current PMA approach is to cluster before setting profiles, the clustering comes essentially for free in step 1. Order the clusters by bit rate from lowest to highest so that the worst profile group is selected for consideration first.
4. Start with the profile with the lowest BR_AP; find the CMs in this profile where BR_AP = BR_PP. If none exists, find the CMs with the lowest absolute difference. Call this CM subset CM_Cluster_i for i=1. Assign to this CM subset CM_Cluster_i a severity measure calculated as the difference between BR_AP and the average BR_PP of all CMs sharing the profile but have a higher BR_PP than the group. So, we are assigning to each CM a cluster number and severity number.
5. Remove this CM_Cluster_1 set of CMs from the set of CMs and recalculate the profiles using PMA (returning to step 1 above).
6. Repeat the steps above to find CM_Cluster_2, and repeat again, until all CMs are clustered into CM_Clusters_i for i = 1 to n for an arbitrary n. Note that n will be larger than the number of profiles. The severity will be used to assign importance for maintenance of the CMs in question, and the CM clusters too.

The work is thus clustered. Sum the severities for the CMs in the cluster. The resulting severities are the priorities to sort by for the work.

Note that the number of profiles has to be discovered from a trusted source such as the engineering team. While the CMTS has a maximum number of profiles possible, operations engineering teams can decide to

¹ In a different version, one could calculate the bit rate for each CM if it could be assigned an optimal profile (max bit rate achievable given how the CM was provisioned) (BP_OP). This could also be defined as the target profile based on provisioning targets.

use a smaller number of profiles when implementing. This setting, therefore, will need to be configured in ProOps or your own PNM application environment for each CMTS.

Measuring the severity of CM clusters in this way presents a tradeoff. Clustering by CMs sharing a profile may correlate with geography but maybe not enough to be efficient with truck rolls. And there is not a guarantee that one cause leads to this CM cluster to share a profile. The theory behind this approach is that one or a few related problems are leading to the lowest profile to bring down a group of CMs. If one or two problems are impacting the group, then this approach should highlight the group of CMs experiencing the problem. But if that is not the case, then the clustering may not work as intended.

Due to the nature of profiles, and the small number of profiles to share among all CMs, one change of a CM's performance can cause re-optimization of the mix of profiles; while some CMs get a better performing profile, others can get a worse one. This makes comparison and prioritization difficult.

If a group of CMs is working fine at a low profile due to being at the end of a line, with low RxMER average values, a profile may be selected in this approach which has no addressable issues. Therefore, human intervention will be needed, and perhaps additional rules or checking is needed to remove false positives.

The need for information from experts to inform the profile management, as well as the fact that this method may indicate false positives, suggests it will be difficult to manage.

This method can easily be augmented with anomaly detection. By adding anomaly detection to the calculation, selecting the worst performing CMs, removing the anomalies, and recalculating the profiles to see the impact, the profile information can be used to prioritize proactive work. Using anomaly detection should reduce the expected false positive rate significantly, and make the method worth considering.

2.2. Prioritize by worst profile modem without anomaly detection

This method is one which we encoded and will provide comparisons later in this document. This procedure below aligns with our measurement labeled Severity(PMA).

The target in this method is the worst performing CMs over all profiles in terms of their impact on the profile. If your PMA calculation uses anomaly detection within it, then anomaly detection is used indirectly but its information is not directly used in this case.

1. For each CM, calculate its profile from PMA based on the number of profiles possible in its CMTS (Assigned Profile, as is typically done with PMA), and the best profile it could have been assigned with unlimited profiles possible (Possible Profile, a profile it would be assigned if it were in a cluster of one). Note this is before consideration of impairments that can be removed.
2. Calculate the bit rates (bit loading) these CMs should achieve with these two profiles (BR_AP, BR_PP).
3. Cluster CMs by profile. Given the current PMA approach is to cluster before setting profiles, the clustering comes essentially for free in step 1. Order the clusters by bit rate from lowest to highest so that the worst profile group is selected for consideration first.
4. Searching over all profiles, find the CMs where $BR_AP = BR_PP$ (plus or minus a small delta). Or, if none exist, find the CM(s) with the smallest difference in $BR_PP - BR_AP$. Call this CM subset CM_Cluster_i for i=1. Assign to CM_Cluster_i a severity measure calculated as the difference between BR_AP and the average BR_PP of all CMs sharing the profile but have a higher BR_PP than the group. So, we are assigning to each CM a cluster number and severity number.

5. Remove this CM_Cluster_1 set of CMs from the set of CMs and recalculate the profiles using PMA (returning to step 1 above).
6. Repeat the steps above to find CM_Cluster_2, and repeat again, until all CMs are clustered into CM_Clusters_i for i = 1 to n for an arbitrary n. Note that n will be larger than the number of profiles. The severity will be used to assign importance for maintenance of the CMs in question, and the CM clusters too.

The work is thus clustered. Sum the severities for the CMs in the cluster. The resulting severities are the priorities to sort by for the proactive work.

Just as in the previous approach, this approach has the same tradeoff due to the clustering utilizing profiles, and using profile impact as the measure of severity.

Like the previous method, this method can easily be augmented with anomaly detection to allow selection by anomaly that potentially would be removed.

2.3. Prioritize by profile bitrate impact using anomaly detection

This method is one which we did not encode for comparison as we expect it to be somewhat like the previous one which we did. However, we think it may be one worth exploring with operators.

The target in this method is the anomalies in the signal, so anomaly detection is explicitly required. But we still use profiles to determine which anomalies to dispatch for.

Assume we have a list of impairments generated for a set of CMs on a given node. Like impairments across CMs, they need to be clustered so that each impairment is associated with a list of CMs that show the impairment, and each CM could be associated with any non-negative number of impairments.

1. For a set of CMs on the same node, and impairments found to appear on the set of CMs, order the impairments found by severity on the CMs and create an association list for each impairment containing the CMs that show the impairment.
2. Find the profiles for each CM on the node using PMA and call these profiles the Assigned Profiles.
3. Take the most severe impairment on the ordered list of impairments, as measured by the amount of impact to RxMER on the subcarriers that it impairs, using a method such as those that follow. Adjust the RxMER per subcarrier for each CM impacted by the impairment as though the impairment was removed (methods like LMS (Cole 1990)), sliding window average or median, EWMA (bi-directional averaged), or FFT-smooth-IFFT methods may be useful here). One reasonable approximation is to assign new RxMER values over affected subcarriers such as the statistical average of the non-affected subcarriers, or a target value for the subcarriers, or for impairments like waves to assign the average of affected subcarriers. Find new profiles for the new set of CMs on the node using the adjusted RxMER per subcarrier.
4. Take the difference in the capacity gains of the newfound profiles and the previously calculated profiles, then assign the delta value as a score to the impairment. Depending on desired approach, cluster the CMs by shared anomaly, profile, or not at all.
5. Update Assigned Profile with the newfound profiles, for all CMs on the node, and repeat the process above (return to step 2) for the next impairment until no more impairments are on the list. All impairments have a severity assigned now.

The default is to cluster by impairment found, as that is given in the process. Alternately, we could sum up severities by CM or profile and assign work based on the severities summed.

By using RxMER impact or bit load as the measure of performance, prioritization is straightforward. There are impacts in these measures due to clustering, as profile changes can have the same impact as discussed above, and removal of an impairment can impact more than one CM. But, the measures of RxMER impact (dB or linear) or bit loading are closely linked to network capacity and service bandwidth, so are excellent candidates for prioritizing work. With AD, a focus on the impairment, and therefore the issue to address, reduces the need to cluster. The clustering only helps to measure the impact of the impairment.

Note: A possible adjustment to this method is to take the square of the difference of RxMER values by subcarrier for each CM and sum those square values so that deviations are measured as a squared loss function instead of linear.

2.4. Prioritize by bit-loading or RxMER per subcarrier impact using anomaly detection

This method is one which we encoded and will provide comparisons later in this document. We label the measurements for this process as Severity(MER) and Severity(Bitload).

The target in this method is the anomalies in the signal again, so therefore anomaly detection is explicitly required here too. But this time we just use RxMER per subcarrier or the bit loading possible, avoiding explicit use of PMA.

Assume we have a list of impairments generated for a set of CMs on a given node. Like impairments across CMs need to be clustered so that each impairment is associated with a list of CMs that show the impairment, and each CM could be associated with any non-negative number of impairments.

1. For a set of CMs on the same node, and impairments found to appear on the set of CMs, order the impairments found by severity on the CMs and create an association list for each impairment containing the CMs that show the impairment.
2. Find the RxMER per subcarrier for each CM on the node and call these the real values.
3. Take the most severe impairment on the ordered list of impairments, as measured by the amount of impact to RxMER on the subcarriers that it impairs, using a method such as those that follow. Adjust the RxMER per subcarrier for each CM impacted by the impairment as though the impairment was removed (methods like LMS, sliding window average or median, exponentially weighted moving average (EWMA) (bi-directional averaged), or FFT-smooth-IFFT methods may be useful here). Note this is a clustering step as well, resulting in CMs that are clustered around an anomaly. One reasonable approximation is to assign new RxMER values over affected subcarriers that are the average of the non-affected subcarriers, or a target value for the subcarriers, or for impairments like waves to assign the average of affected subcarriers. Call this adjusted RxMER per subcarrier the adjusted values. Apply this step to all impacted CMs in the group.
4. Sum the improvement in bit rate (best possible bit loading) that is gained by removal of the impairment, over all impacted CMs, and assign this result as a score to the impairment. Alternately, one can use the improvement in RxMER per subcarrier values summed over subcarriers.
5. Update the real RxMER values with the adjusted RxMER values for all subcarriers, for all CMs on the node, and repeat the process above (return to step 3 with the updated RxMER per

subcarrier values) for the next impairment until no more impairments are on the list. All impairments have a severity assigned now.

The CMs are clustered by anomaly, so that can be used easily by default. Individual CMs can drive the work too, depending on how the user wants to apply the severity and sort for work.

As this method uses AD and measures the impact of the removal of the anomaly, the advantages of the previous method apply here as well.

Note: A possible adjustment to this method is to take the square of the difference of RxMER values by subcarrier for each CM and sum those square values so that deviations are measured as a squared loss function instead of linear.

2.5. Prioritize by bit-loading or RxMER per subcarrier without anomaly detection

This method targets poor performing CMs without the use of AD or even PMA or profile information of any kind. It instead goes to the source of profile information, the downstream RxMER per subcarrier values for the CMs involved. We encoded this measurement for comparison later, indicating the measurements as NormalScore(MER) and NormalScore(Bitload).

This simple approach is offered as a method to begin using and developing from. The PMA application is shared in limited ways with vendors, and the CableLabs AD is shared with vendors as an executable only. Not every operator will implement PMA. Such operators are able to make use of this method here too, but some of the methods using PMA and AD are expected to perform better for many operator applications.

1. For a set of CMs on the same node, find the RxMER per subcarrier for each CM on the node, and call these the real values.
2. For each CM, adjust the RxMER per subcarrier for each CM through smoothing (methods like LMS, sliding window average or median, EWMA (bi-directional averaged), or FFT-smooth-IFFT methods may be useful here). One reasonable approximation is to assign new RxMER values over affected subcarriers that are average of the non-affected subcarriers, or a target value for the subcarriers, or for impairments like waves to assign the average of affected subcarriers. Call this adjusted RxMER per subcarrier the adjusted values.
3. Take the sum of the squares or absolute value of the differences in each subcarrier between the real and adjusted RxMER values and call this the severity. A bit loading calculation can be done here as well, and the overall bit loading may be used as the measure instead (inverse of bit load, or difference in bit loading from the real and smoothed RxMER per subcarrier results).
4. Prioritize and sort the CMs by their severities.

GIS information would be needed to cluster work without alternatives such as by anomaly. CM performance may inform that clustering too, as experts have suggested.

3. Details, Discussion, and Evaluation

3.1. A note on using these results in a PNM framework

To use these approaches as workers in ProOps, and to have a solution that turns data into actionable information, we need to have elements that collect data and calculate statistics (observe), assign priorities and cluster work (orient), sort the clusters and choose which are opportunities (decide) on which to act (act). The elements in these flows sometimes do not follow an order for OODA, but the elements are

represented. Some of the elements which may be separate workers when coded for ProOps, and the roles they may provide, are described next.

- Observe
 - RxMER per subcarrier
 - AD
 - Profile
- Orient
 - RxMER improvement or severity
 - AD severity
 - Profile improvement
 - Anomaly matching
- Decide
 - Severity assignment
 - Sorting results
- Act
 - Dashboard
 - Work list

For an initial configuration, here are the general configuration elements and steps.

- Collect RxMER per subcarrier every 4 hours, and profiles if relevant, then immediately run through any AD, or statistics calculators.
- Calculate any severities or scores needed.
- If the severity calculation did not require clustering yet, then execute on any clustering needed.
- Execute next on any added data to collect for future AD or impairment detection training, etc.
- Calculate any statistics needed based on the clustering, including calculating the cluster severities.
- Sort the clusters or single entities by severity, and calculate any needed data for the dashboard and prioritized work list, which may have a threshold value applied as criteria for inclusion in the work list.
- Present results in a desired file format or data base, along with updating the dashboard.

3.2. A note on clustering work by anomalies

As we examine anomalies found in CMs, we will calculate their impact on RxMER per subcarrier or on profiles and assign a value to the CMs where the impairment appears. But in doing this, we want to also combine CMs by impairment in many cases. Combining by profile is done in some methods by default. But working by anomalies may not result in clustered CMs unless the anomalies are compared. An untested approach we suggest is as follows.

1. Cluster anomalies by their features:
 - a. For waves, match by cavity wall size first (distance between peaks or troughs), then severity.
 - b. For slope, by slope value.
 - c. For other impairments, by subcarrier frequency or data ranges.
2. For each cluster, add the net severity scores of the MAC addresses in the clusters; call this the cluster severity.

Several methods may be useful to consider here. For waves, normalize the slope to 0, then take an IFFT to find the frequency values most prominent, and finally correlate the transformed data. Likewise, correlate the untransformed data. Methods to correlate these data include k-means, difference tests like paired-t and Chi-Squared, and some statistical parameter tests such as moment tests.

There are several challenges with comparing anomalies to cluster work, and we leave this work for future consideration. We may start with investigating the IFFT approach with checking for correlation of the transformed and untransformed tests for waves, and use difference tests for frequency bound anomalies (matched by anomaly type or not).

3.3. What to monitor

Several statistics may be worth monitoring including the following.

- Table of CMs by severity, ordered by severity
- Table of CMs by hours in the top worst (configurable) by net severity
- Plot of clusters by number of CMs in the cluster and severity
- Plot of the top worst (configurable) CMs or clusters by severity, over time
- Overall node health score by severity total

3.4. Training data for anomaly detection

Training on new data of the AD needs a human supervisor to be involved to label samples or make corrections. Labeling samples means manually adding labels to the new samples. We can also use the AD to perform detection on the new samples, and the human supervisor makes corrections to the localization and classification results. Even if the prediction seems precise enough, small corrections are still necessary to make the labels even more precise. This work is needed because the feed forward neural networks are deterministic in terms of making predictions on the same input data. If we use the same data with the label that's generated by the AD itself for training, it's only reinforcing the model to what it has learned and not helping the model to learn from the new samples.

Each type of spectrum data will use a separate training process to train a separate model using labeled datasets, because different spectrum data may reflect impairments differently. Also using separate models can help stabilize the model and improve efficiency and accuracy.

3.5. Bit loading

Bit loading is handled differently depending on whether the solution is based on PMA or AD (including non-AD RxMER calculations).

With PMA, the bit loading gain is represented by the J value. PMA calculates profiles, then calculates a relative capacity gain by using dynamic profiles over 256-QAM flat profiles as the J value. The J value can be used to find an overall channel capacity gain in bit load. By calculating a set of modulation profiles, making an adjustment to the data, then calculating a new set of profiles, the differences in total bit loading can be compared. But realize that by looking at profiles and making changes based on those, the gain in bit loading depends on the other CMs on the impacted channels, so making a change to a CM or one or more channels impacts multiple CMs, and potentially more than one profile, so an overall impact has to be considered.

With AD, anomalies are first detected by the AD. As we know their type and frequency range from the AD, we first calculate the average bit loading from the average MER value from all subcarriers, then calculate the difference between subcarrier (within the anomaly range) bit loading values and the average bit loading. So, for each anomaly we have a total bit loading loss calculated. These scores can be weighted by anomaly types and added together. This is how we calculate the bit loading in these approaches using AD; RxMER-based bit loading calculations are similar.

3.6. Measuring performance

Because these methods have a few very different methods for measuring severity, it is difficult to compare the performance of some of these approaches. Likewise, combining these methods is possible but a weighting scheme for that might need to be developed. Further, there really are two ways to think about the operational performance of these approaches: the alignment of severity to actual problems in service and plant, and the alignment of severity to financially responsible repair work. The methods informing these severity measures are separately comparable by their ability to collect valuable information to support manual troubleshooting and repair. In addition, the calculation time required (performance) of the methods is a consideration too.

For the methods we encoded, we can provide computational performance comparisons. But without field studies to compare effectiveness, we can't provide any useful guidance on the performance of one approach over another in terms of the effectiveness of the work performed as a result. We can, however, provide a comparison of the information that each approach provides to the technicians, and our risk assessment of which methods should most closely align with effective work. We leave operational performance comparisons to the operators, and we look forward to working with any operator willing to work with us to optimize these methods.

3.7. Comparison of measurements

We applied several of our methods to a large data set of RxMER per subcarrier values and calculated the output measurements of each. Where we could apply different measurements to the same CM, we did so; the results are plotted in the figures that follow.

In these figures, when referring to bit load (labeled Bitload) and RxMER (labeled as MER), severity score and normal score are different ways to calculate similar statistics, and thus potentially used in the same way.

- **Severity score** is the sum of separate anomaly scores (calculated against the sample average in this case), so these two severity scores are from the AD involved methods.
- Alternately, the **normal score** calculation doesn't include AD. It only calculates the "area" below the sliding median across the sample. These normal scores are therefore a measure of variability. We can use them as a severity measure like severity score.

However, both severity and normal scores are from the perspective of the single CM.

Alternately, the channel score is the sum of all CM scores on the same channel; the optimal total capacity of the channel with a certain number of profiles and other constraints can be estimated using profiles generated by PMA, so is the important measure for PMA consideration. In the figures below, we refer to channel score as Severity(PMA).

Note that Severity(Bitload) would be equivalent to Severity(PMA) if there were unlimited profiles (or as many or more profiles as CMs, so that each CM can have its perfectly matched profile).

See Figure 1 for the relationship between Severity(MER) and NormalScore(MER). Note all scores are positive, and we see a cone shape relationship. While the two approaches correlate, this graph does suggest that one measure will prioritize work differently than the other. The top few for one measure are different than the top few for the other.




Figure 1 - Severity versus normal score from RxMER

Figure 2 shows the same comparison between severity and normal score, but for bit loading. Again, all measurements are positive. But there are several observations that score differently on one measure with little change in the other. Again, taking the top opportunities with one measure and the top for the other measure may not have any overlap, or very few. The two approaches will suggest different work to consider.




Figure 2 - Severity versus normal score from bit load

Figure 3 compares the severity measures of MER and bit load. Notice a high degree of correlation. This graph suggests that using either measurement is sufficient. And due to the way that bit loading is calculated, this relationship is expected.




Figure 3 - Severity from RxMER value versus severity from bit loading value

With profile management added to consideration, we can calculate the severity of the PMA measurement and compare it to the other severity measurements. Figure 4 shows the MER severity compared to the PMA severity. As expected, there are clusters where addressing the problem may have negative impact on the overall profile, and those do have smaller MER severity. Certainly, if removing an impairment will result in a net overall lower performance from PMA, that might not be work you want to prioritize (there are other conditions that may change that decision, and there may be different approaches to PMA that avoid the conflict, as we discuss later).

However, note that at least for this set of data, those with negative impact after considering profiles are less severe. Therefore, if we use RxMER per subcarrier to prioritize PNM work, the risk may be low of having a negative or small impact after considering profiles. It may be sufficient just to check the profile impact when in doubt, if using RxMER per subcarrier as the basis for selecting proactive work.




Figure 4 - Severity from RxMER value versus severity from PMA value

Figure 5 compares severity of bit loading and PMA. Like Figure 4, we see the same relationship and the conclusions are the same, as expected.




Figure 5 - Severity from bit loading versus severity from PMA value

Figure 6 relates the normal scores of MER and bit loading. Like Figure 3, we see a strong correlation. But this time there is a large number of high normal scores for bit loading that have small MER normal scores. This suggests that, without the ability to find impairments in an automated fashion, there is good reason to look at both measurements and perhaps look deeper to determine which is best. But this is also another reason for using an AD.




Figure 6 - Normal score from RxMER value versus normal score from bit loading value

Figure 4 and Figure 5 suggest looking at severity for MER and bit loading for only positive PMA severity. Figure 7 shows the conditional relationship, and it strongly resembles Figure 3. This result suggests that the dynamics of PMA cannot be well predicted by just looking at improvements in MER or bit loading, so some consideration of profile impact is warranted for selecting PNM work.




Figure 7 - Severity for RxMER value versus bit loading for CM that score a positive PMA severity

Next, we want to look at individual anomaly types to see if some types of anomalies exhibit different relationships. If we find that some anomaly types appear to score generally higher, or have different relationships with profiles, then we may want to treat them differently. But as we see in the figures that follow, there seems to be no difference in these relationships by anomaly type.

Figure 8 through Figure 13 are repeats of Figure 1 through Figure 6 respectively, but for only LTE anomalies. Note the patterns are respectively very similar, leading to the same conclusions for just LTE anomalies.




Figure 8 - Severity versus normal score from RxMER for LTE ingress impairments




Figure 9 - Severity versus normal score from bit load for LTE ingress impairments




Figure 10 - Severity from RxMER value versus severity from bit loading value for LTE ingress impairments




Figure 11 - Severity from RxMER value versus severity from PMA value for LTE ingress impairments




Figure 12 - Severity from bit loading versus severity from PMA value for LTE ingress impairments




Figure 13 - Normal score from RxMER value versus normal score from bit loading value for LTE ingress impairments

With five separate impairment types to identify, and several modems showing multiple impairment types, showing all graph combinations would fill several pages with little value. Therefore, we are showing a few sample graphs for other single impairments, and for those modems with multiple impairments in Figure 14 through Figure 18, which follow.




Figure 14 - Severity from RxMER value versus severity from bit loading value for wave impairments




Figure 15 - Severity from RxMER value versus severity from bit loading value for rolloff impairments




Figure 16 - Severity from RxMER value versus severity from bit loading value for suckout impairments




Figure 17 - Severity from RxMER value versus severity from bit loading value for spike impairments




Figure 18 - Severity from RxMER value versus severity from bit loading value for modems reporting multiple impairment types

These plots suggest a few important results which inform maintenance decisions.

A reliable method for identifying anomalies in an automated way is important. Without analyzing the tradeoffs between statistical methods and machine learning techniques for this problem, we only suggest operators use anomalies to identify the opportunities to address because they are indications of cable plant problems.

Further, we suggest using either RxMER per subcarrier or bit loading as a measurement basis, with a measurement of the difference in either as a way to prioritize work. Equally, a plant quality measurement could be used if it is a comparison between the actual performance and intended performance.

Profile management must be considered when selecting proactive maintenance work. Looking at the provided figures, it is conceivable that only the most impactful impairments are addressed, and those appear to have positive impact on overall profiles. But not all PNM activities will yield positive net profile impact, so some consideration should be given here. We explore this issue a bit more next.

3.8. Manual simulation of proactive maintenance impacts

As an exercise to further demonstrate these relationships, we next conduct a manual simulation of maintenance performed on a set of modems. We select a sample of modem data from the larger data set which generated the figures previously discussed. We then select the most significant impairment on a measure of performance, remove the identified impairment from the RxMER per subcarrier data of all modems where it appears by a mathematical adjustment, recalculate profiles and measurements, then repeat a few times. We then report the results of this experiment in a few forms.

While many of the methods we outline in this paper are CM or CM cluster centric, this approach demonstrates using two of the methods as a measurement but in an impairment centric maintenance approach, which we believe is an operationally effective way to perform PNM. The impact on profiles however is examined here too, as we can already see it should be considered for an overall network service impact perspective.

In the comparison done here, the method of section 2.2 is augmented to be an anomaly driven version of using Severity(Profile) as the measure, and this is compared to the method of section 2.4 where we prioritize by bit load impact from correcting the impairment (Severity(Bitload)). Therefore, looking at the difference should show something about the impact of profiles on the simulations.

We make one large assumption in this analysis worth mentioning: we are identifying anomalies in multiple CMs and correlating anomalies so that we presumably remove one anomaly from multiple CMs where it appears. Therefore, the assumption is that the anomaly is indeed one anomaly appearing in multiple CMs, whereas it may not be the same anomaly but rather more than one that coincidentally appear the same in the RxMER per subcarrier data. While the strength of this assumption is related to our ability to accurately cluster by anomaly, the use is comparative in this simulation so should not bias the results greatly. The intent is to see the impact of profiles on PNM selection, and it should still show that impact.

For this simulation comparison, we selected a set of 526 CMs on the same section of plant. Figure 19 shows the RxMER per subcarrier data of all these CMs overlaid on the same plot. Notice related impairments are indicated.




Figure 19 - 526 CMs on a common channel, many experiencing impairments that appear related

The simulation results are shown in Figure 20 and Figure 21, as well as Table 1 and Table 2. Figure 20 is a graph of the total bit load for the set of CMs at each iteration, where an iteration is the mathematical removal of an impairment from the RxMER values, with the impairment selected by the two compared methods: Severity(Bitload) and Severity(Profile). Note that for the first few impairments, there is little difference in the methods. In fact, seeing the detailed results in the two tables, the first three impairments selected by the two methods are the same. After about seven iterations, there is little difference in the two approaches. These results suggest that either approach may be sufficient for selecting PNM work. But more data and simulations would be needed to prove this hypothesis.

Looking at Figure 21, we see that the profile management results, as indicated by J value, vary more. Thus, at least for the selected PMA used here, there is impact on PNM selection and the impact of PNM on service. Of course, this impact is not likely detectable by customers or even applications in many cases, but it may be in extreme cases.

This result suggests that the profile management approach implemented should be made robust to PNM activities. Fortunately, this is easy: if each CM's RxMER is the same or improved after PNM, and a recalculation of profiles yields a solution that is further from the objective, then keep the existing profile solution set, and consider recalculating until a better solution is found. As many PMA are heuristics, you can always change the order of inputs and get a different answer. As profile calculation is fast, and profiles are often kept for long periods of time, there is little reason not to calculate several solutions and select the best according to operations and business objectives.




Figure 20 - Cumulative bit load improvements over multiple rounds of impairment removal for both methods




Figure 21 - Cumulative J value (PMA) improvements over multiple rounds of impairment removal for both methods

Table 1- Results prioritized by bit load (Severity(Bitload))

Round	Anomaly	Location (Subcarrier ID)	Num CMs	Severity (Profile)	Severity (Bit load)	Num CMs (Better Profile)	Num CMs (Worse Profile)
0	LTE	start: 1128, end: 1281	39/526	0.696	3356	352	174
1	LTE	start: 1010, end: 1097	90/526	0.156	3157	193	333
2	WAVE	start: 0, end: 1399	40/526	0.536	2828	178	323
3	LTE	start: 1346, end: 1376	23/526	-0.013	951	229	272
4	ROLLOFF	start: 0, end: 138	9/526	-0.001	923	357	163
5	SUCKOUT	start: 992, end: 1114	27/526	0.564	814	433	77
6	ROLLOFF	start: 1330, end: 1399	3/526	-0.192	182	220	109
7	SUCKOUT	start: 1082, end: 1318	8/526	0.158	142	90	219
8	LTE	start: 1001, end: 1028	31/526	-0.138	75	241	127
9	LTE	start: 1259, end: 1288	2/526	-0.130	33	98	78
10	SPIKE	start: 16, end: 23	2/526	0.000	2	0	0
11	SPIKE	start: 1348, end: 1354	2/526	0.000	0	0	0
12	SPIKE	start: 1108, end: 1114	3/526	0.000	0	0	0
13	SPIKE	start: 1227, end: 1234	2/526	0.003	0	2	51
14	SPIKE	start: 712, end: 718	2/526	0.000	0	0	0

Table 2 - Results prioritized by J value (equivalent to Severity(PMA)).

Round	Anomaly	Location (Subcarrier ID)	Num CMs	Severity (Profile)	Severity (Bit load)	Num CMs (Better Profile)	Num CMs (Worse Profile)
0	LTE	start: 1128, end: 1281	39/526	0.696	3356	352	174
1	LTE	start: 1010, end: 1097	90/526	0.156	3157	193	333
2	WAVE	start: 0, end: 1399	40/526	0.536	2828	178	323

Round	Anomaly	Location (Subcarrier ID)	Num CMs	Severity (Profile)	Severity (Bit load)	Num CMs (Better Profile)	Num CMs (Worse Profile)
3	SUCKOUT	start: 1082, end: 1318	8/526	0.062	146	79	250
4	SUCKOUT	start: 992, end: 1114	27/526	0.159	810	335	178
5	ROLLOFF	start: 0, end: 138	9/526	0.193	923	128	125
6	LTE	start: 1259, end: 1288	2/526	0.053	33	106	23
7	LTE	start: 1346, end: 1376	23/526	0.122	951	47	305
8	SPIKE	start: 712, end: 718	2/526	0.000	0	0	0
9	SPIKE	start: 16, end: 23	2/526	0.000	2	0	0
10	SPIKE	start: 1227, end: 1234	2/526	0.000	0	0	0
11	SPIKE	start: 1108, end: 1114	3/526	0.000	0	0	0
12	SPIKE	start: 1348, end: 1354	2/526	0.000	0	0	0
13	ROLLOFF	start: 1330, end: 1399	3/526	-0.005	182	147	40
14	LTE	start: 1001, end: 1028	31/526	-0.332	75	160	154

4. Conclusion and Future Work

Profile management has a clear influence on the impact to service that results from PNM. Therefore, some consideration of profiles is important when selecting proactive work to do. It is advisable that operators at least look at the impact on profiles when determining whether to conduct PNM, and it is likely a factor in prioritizing proactive work.

Anomaly-driven PNM, finding and eliminating anomalies, appears to be the most popular approach that operators take for their PNM; but modem-driven PNM is an option to consider for some, as the data that most informs PNM comes from the CMs. Likewise, profile-driven PNM is an option to consider as well. But ultimately, each of these factors needs to be considered in some way in deciding which PNM work to take on.

Though this paper does not compare all options, an examination of the options presented here does suggest that impairment-driven identification of maintenance opportunities is a very good approach, that problem severity can be measured by impact on CMs, and that it is necessary and complimentary to PNM decision making to consider the impact on profiles when determining the capacity impact on the network for selecting maintenance. Said differently—find anomalies, investigate their impact on customers’

service, estimate the impact due to profiles and on network impact expected from the work, and use these pieces of information to prioritize and select PNM opportunities. It may at times be best to let profiles drive PNM maintenance decisions, in fact.

Seeing that some network repairs can have negative impact on profiles, it suggests more work on profiles is warranted too. It may be simple enough to set a rule that says, if all CMs improve, then the recalculation of profiles can only result in a net benefit, or a benefit to each CM, or another appropriate rule, before consideration of replacing profiles. In other words, do not harm first. Alternately, it may be useful to develop algorithms that set profiles to be robust against these changes. Likewise, it may be worth investigating different objectives when setting profiles altogether.

Further computer simulations of various PNM selection criteria are warranted by this work. The two limited simulations we ran suggest that very significant PNM opportunities can have significant impact on network capacity, whereas smaller ones may not. Further, there are several measures that can be used to select PNM opportunities, and they may all be sufficient for many operations, though that should be checked with further simulation and field experiments. The limited simulations we ran also suggest that profiles recalculated after plant improvements may not always improve capacity for everyone or overall, but that can be easily addressed in software or in operations decisions.

At least, consider this work as a roadmap for analyzing PNM strategies for operations. The graphical comparisons are simple to do with your own data, and running simulations of various approaches can be very informative. Extending this work into cost analysis may be the necessary step to procure funding for operational improvements.

Our best hope is to work with operators to test some of these approaches and determine which are the most cost effective to operations and impactful to service overall. Selecting proactive work requires identifying the opportunity, measuring it, estimating the benefit of addressing it, and then executing on those opportunities with the most promise. Therefore, the best measure of success is to determine the cost and benefit of the work being done, then adjust the opportunity selection process to optimize the program on those measures. The field work that has to be done is to determine which of the approaches outlined in this paper, or alternative approaches, are best suited to optimize the maintenance program. Looking at the conclusions that follow, we already have a good idea where to start. Fortunately, for well-equipped or unequipped operators alike, CableLabs has ProOps encoded and ready to assist.

Certainly, there remains a lot of confirmation work with operators before all this can be applied optimally to operations. While identifying the anomalies and measuring their significance to service are important steps, the PNM work to remove the anomalies has to be done to measure cost and impact, so that both can be modeled better for more accurate and operator-specific decision making.

Abbreviations

AD	anomaly detection
CM	cable modem
CMTS	cable modem termination system
dB	decibels
DOCSIS	data over cable service interface specification
IFFT	inverse fast Fourier transform
ISBE	International Society of Broadband Experts

LMS	least mean square
LTE	long term evolution
MER	modulation error ratio
OODA	observe, orient, decide, act
PMA	profile management application
PNM	proactive network maintenance
ProOps	proactive operations (platform)
QAM	quadrature amplitude modulation
RxMER	receive modulation error ratio
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

[1] Cole, T.J., “The LMS method for constructing normalized growth standards,” European Journal of Clinical Nutrition 44(1): 45-60, 1990

[2] Jingjie Zhu, Karthik Sundaresan, Jason Rupe, “Proactive network maintenance using fast, accurate anomaly localization and classification in 1-D data series,” proceedings of the International Conference on Prognostics and Health Management (ICPHM), June 2020, 978-1-7281-6286-7.

[3] Ron Hranac, James Medlock, Bruce Currivan, Roger Fish, Tom Kolze, Jason Rupe, Tom Williams, Larry Wolcott, “Characterizing network problems using DOCSIS(R) 3.1 OFDM RxMER per subcarrier data,” 2019 SCTE-ISBE and NCTA.

[4] Greg White, Karthik Sundaresan, “DOCSIS 3.1 Profile Management Application and Algorithms,” 2016 SCTE-ISBE and NCTA.

[5] Jason Rupe, Jingjie Zhu, “Kickstarting proactive network maintenance with the proactive operations platform and example application,” 2019 SCTE-ISBE and NCTA.

[6] Jason Rupe, “A general-purpose operations cost model to support proactive network maintenance and more,” 2019 SCTE-ISBE and NCTA.

Streaming Telemetry Data from the Home Network using OpenWrt Access CPE

A Technical Paper prepared for SCTE•ISBE by

Shlomo Ovadia, Ph.D.
Senior Principal Engineer
Charter Communications
14810 Grasslands Drive, Englewood, CO 80112
720-536-1686
Shlomo.ovadia@charter.com

Deependra Rawat
Sr. Software Developer
Charter Communications
14810 Grasslands Drive, Englewood, CO 80112
c-deependra.rawat@charter.com

Dan Lynch
Sr. Software Developer
Charter Communications
14810 Grasslands Drive, Englewood, CO 80112
c-daniel.lynch1@charter.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Home Network Architecture.....	3
3. Access CPE OpenWrt Software Architecture.....	4
4. Access CPE OpenSync™ Software Architecture.....	7
4.1. iPerf Speed Test.....	8
4.1.1. STeMTA Plugin:.....	8
4.1.2. STLNeMTA Plugin:.....	9
4.1.3. Cloud Security	9
5. Streaming Telemetry Data Path	9
6. Grafana Dashboard Design.....	11
7. Comparison with Other Streaming Telemetry Methods	15
8. Conclusion	16
Acknowledgment	17
Abbreviations.....	17
Bibliography & References	18

List of Figures

Title	Page Number
Figure 1: Current and Proposed Home Network Architecture Diagrams (Configurations A and B).....	4
Figure 2: OpenWrt Integrated with OpenSync™ Software Architecture	5
Figure 3: OpenSync™ Software Architecture with the Connectivity to the OpenSync™ Cloud.....	9
Figure 4: Telemetry Data Path from the Access CPE Device to the Grafana Dashboard	10
Figure 5: Hierarchical Color-Coded Grafana Dashboard with the Key Telemetry Components.....	12
Figure 6: Reported Status Level 2 Access CPE Metrics in the Last Hour.....	13
Figure 7: CPU Utilization of Access CPE vs. Time (Level 3) Reported in the Last Hour	13
Figure 8: Home Network Traffic Parameters for All the Wirelessly Connected Client in the Home Network	14
Figure 9: Downstream DOCSIS Channel Information Status	15

List of Tables

Title	Page Number
Table 1: Summary of OpenSync™ Managers' Functionality and Status	7
Table 2: Abbreviations Table	17

1. Introduction

In the last decade, home network architectures have become more complicated due to advances in wireless technology, and the explosion of different types of wirelessly connected devices such as Internet-of-Thing (IoT) devices, cell phones, tablets, laptops, gaming devices, etc. In fact, the U.S. smart home market is expected to show a compound annual growth rate of 16.9%, reaching 46.6B by 2024 [1]. In this home network architecture, some cable Multiple System Operators (MSOs) deploy a two-box solution for both residential and Small and Medium-size Business (SMB) customers where one box is an access CPE device (i.e., cable modem or an ONU), and the second box is a wireless router. Cable operators typically have limited information about the access CPE device's health status, and no information about the customer's home network, including what type of clients are connected to the home network, and their bandwidth usage vs. time. Such home network health and traffic information would be very useful to the Cable operators in order to enhance their customers' experience by optimizing the customer's home network traffic, prevent potential field issues, and be able to introduce new services such as customer's home network management and security. Furthermore, some Cable operators have already begun the transition towards cloud-based management of wireless routers and other devices, which is not supported by the currently field-deployed access CPE devices.

In this paper, the challenges with the current home network architecture are first explained. Then, the proposed home network architecture with an agile software stack on the access CPE device is discussed. This includes the agile software stack and components with the cloud-based management of the access CPE device to enable streaming of all the telemetry data. Third, the streaming telemetry data path, including the components and operation of the cable MSO's Streaming and Analytics platform are explained. Fourth, the organized hierarchical Grafana dashboard design with all of the different types of streamed data telemetry, including home network traffic metrics, D3.1 eMTA health metrics, DOCSIS RF info, event alarm and notifications is explained. A comparison between the OpenWrt-based streaming telemetry method in the paper and other streaming telemetry methods is then reviewed. Finally, the benefits to Cable operators using the OpenWrt-based streaming data telemetry method for access CPE devices are summarized.

2. Home Network Architecture

Cable MSOs' common customer home network is a two-box solution, consisting of a D3.1 eMTA device, which is connected to an HFC network via a coaxial cable on the WAN port, and connected to a Wi-Fi router on the LAN port via an Ethernet cable as shown in Figure 1 configuration A. For a Fiber-To-The-Home (FTTH) deployment, the access CPE device is an Optical Network Unit (ONU) connected to a Passive Optical Network (PON). A Wi-Fi router with a number of Wi-Fi APs or multiple Wi-Fi routers are typically used in larger homes for wireless mesh connectivity, resulting in improved data throughputs for all of the wirelessly connected devices in the home network. The software stack on the D3.1 eMTA device is monolithic, customized for each Silicon and OEM vendor, and does not provide any metrics about the health of the access CPE device and the wirelessly connected clients. The D3.1 eMTA monolithic software stack results in long and costly validation testing cycles before the new firmware that meets the Cable operator's requirements can be deployed. Furthermore, Cable operators are moving toward a cloud-based infrastructure for both command/control and streaming telemetry, which the current access CPE device's monolithic software architecture does not accommodate.




Figure 1: Current and Proposed Home Network Architecture Diagrams (Configurations A and B)

3. Access CPE OpenWrt Software Architecture

OpenWrt is an open-source project for embedded Operating System (OS) based on Linux. It was selected since it is highly-flexible open-source OS with a large ecosystem of vendors and developers that enable cable MSOs to rapidly develop new features and plugins that can also be containerized [2]. One of the key built-in benefits of OpenWrt OS is a full carrier-grade IPv4/IPv6 routing functionality on the access CPE device with no need to redesign the hardware. Moving the routing functionality from the connected Wi-Fi router to the access CPE device enables cost optimization by using a generic Wi-Fi AP as the second box, and focusing the AP performance on the wireless connectivity in the home network.

To address the challenge explained above, an agile OpenWrt-based software stack was developed as a Proof of Concept (PoC) using D3.1 eMTA device operating in a home network architecture as shown in Figure 1 configuration B. The agile OpenWrt software stack is integrated with an OpenSync™ layer, a Silicon vendor Software Development Kit (SDK), and the Message Queue Telemetry Transport (MQTT) server architecture as shown in Figure 2. In addition, the CM and voice firmware was loaded on the access CPE device.

MQTT is a lightweight Machine to-Machine (M2M) transport communications protocol [3]. The D3.1 eMTA streams the telemetry data statistics to the Grafana dashboard via the MQTT server that is hosted in the OpenSync™ cloud, which forwards the data to the cable MSO's Streaming and Analytics platform. MQTT supports various authentications and data security mechanisms (using a script to generate security certificates). The Grafana tool was selected since it is a multi-platform open-source analytics and interactive visualization web application that users can customize to create complex monitoring dashboards [4].

The OpenSync™ cloud is composed of a Network Operations Center (NOC) and OpenSync™ controller for managing a network of OpenSync™-enabled devices. Cable operators can establish their own

OpenSync™ cloud by obtaining an OpenSync™ source-code license. The OpenSync™ cloud provides operator-friendly services, including:

- Device and firmware management
- Inventory and billing system
- Network performance control
- Onboarding and provisioning of field-deployed devices
- Telemetry reporting and data analytics
- Network operations, and customer support.

The NOC in the OpenSync™ cloud translates and communicates management commands in a single Wi-Fi AP network and in mesh multi-AP Wi-Fi network via Open vSwitch Database (OVSDB) distributed database commands. The OpenSync™ controller utilizes the OVS implementation and OpenFlow protocol for networking, and MQTT protocol for receiving state and telemetry data from OpenSync™-enabled devices. Specifically, the OpenSync™ controller provides the necessary command and control services such as:

- Network Status
- IP Address (displayed)
- Network Mask
- DHCP Status
- Parental Control
- Speed Test initiation and results
- Reset and Reboot device

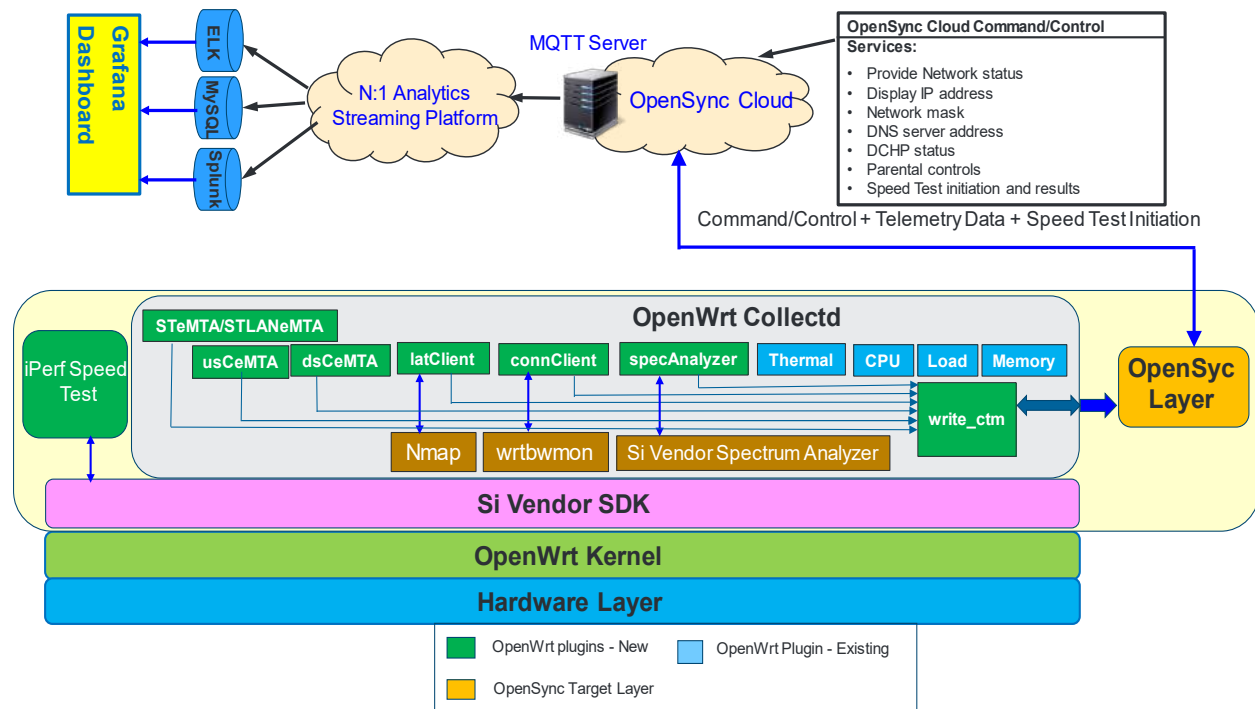


Figure 2: OpenWrt Integrated with OpenSync™ Software Architecture

In this software architecture a smart remote agent based on an OpenWrt data collector open-source software component called `collected` was used as shown in Figure 2 [5]. The `collected` component gathers metrics from various sources, e.g. the operating system, applications, log-files and external devices, and stores this information or makes it available over the network. Those statistics can be used to monitor systems, find performance bottlenecks (i.e. *performance analysis*) and predict future system load (i.e. *capacity planning*). The `collected` component, which offers a variety of Plugins (software programs), is used to collect different types of telemetry data from a few Wi-Fi routers. The `collected` component's default reporting time interval is 30 seconds, but other configurable time intervals can be selected. New capabilities and functionality (green-colored boxes in Figure 2) were added to the `collected` software components. For example, the smart remote agent can be used to run a specific test such as measure the IPv4/IPv6 DOCSIS round-trip latency as explained below using the `eMTALat` plugin, read the collected measurement data, and stream the data to the service operator's streaming and analytics platform. The blue-colored boxes are existing supported OpenWrt `Collected` plugins that are utilized in this architecture. The `collected` software components are integrated with the Silicon vendor's SDK and with the supported OpenWrt OS. It should be pointed out that other custom plugins and shell scripts can be developed, and the listed plugins below are just a sample framework.

Green-coded `collected` plugins descriptions:

1. **usCeMTA:** Software plugin to pull all the DOCSIS upstream channel information used by the D3.1 eMTA (RF level, channel frequency, etc.).
2. **dsCeMTA:** Software plugin to pull all the DOCSIS downstream channel information used by the D3.1 eMTA (RF level, channel frequency, etc.).
3. **latClient:** Software plugin that measures and reports the round-trip latency from the D3.1 eMTA to each of the wirelessly connected devices in the home network based on their IP address or MAC address as shown in Figure 1.
4. **connClient:** Software plugin that measures and reports the number of transmitted and received packets from each of the wirelessly connected devices in the home network.
5. **specAnalyzer:** Software plugin to obtain the RF downstream and upstream spectrum of the Access CPE device.
6. **eMTALat:** Software plugin that measures and reports the minimum, maximum, and average round-trip DOCSIS latency between the D3.1 eMTA and the connected CMTS. First, it initiates a trace-route command to get CMTS IPv4 and IPv6 addresses. Then, it starts ICMP request and reply commands to measure the DOCSIS latency between CM and CMTS and stores the test results in separate files for IPv4 and IPv6. The `eMTALat` plugin reads results from these files and send them to `write_ctm` plugin, which in turn sends the measured DOCSIS latency results to the OpenSync™ layer's SM (not shown in Figure 1).

Wrtbwmon:

`wrtbwmon` is a small and basic shell script designed to run on Linux powered routers (OpenWRT, DD-WRT, Tomato, and other routers where shell access is available). It provides per user bandwidth monitoring capabilities and generates usage reports [6].

Nmap:

Network Mapper or Nmap is a free and open source utility for network discovery and security auditing [7]. Many systems and network administrators also find it useful for tasks such as network inventory, managing service upgrade schedules, and monitoring host or service uptime. Nmap uses raw IP packets in novel ways to determine what hosts are available on the network, what services (application name and version) those hosts are offering, what operating systems (and OS versions) they are running, what type of packet filters/firewalls are in use, and dozens of other characteristics (nmap.org).

write_ctm collectd Plugin:

The **collectd** service calls its collection plugins over a configurable time period. The collectd daemon collects various statistics from the device for aggregation and forwarding to a desired destination. The write_ctm plugin aggregates the stats, converts them to Protobuf format, and sends the collected telemetry data to the OpenSync™ STATS Manager (SM), which in turns forwards the collected telemetry data to the OpenSync™ Queue Manager (QM). The OpenSync™ QM sends the collected telemetry data to the cable MSO's Streaming and Analytics platform via the MQTT server. It should be pointed out that collectd plugin can also send string-formatted event notification based on system defined threshold levels. The write_ctm plugin registers its write API to collectd, and fetches data from the existing plugins on the expiration of every time period. No changes are needed to the existing plugins.

The blue-coded collectd plugins (Thermal, CPU, Load, Memory) are standard open-source collectd plugins that were added to the list of collectd plugins as shown in Figure 2.

4. Access CPE OpenSync™ Software Architecture

Figure 3 shows the key OpenSync™ software architecture components integrated with OpenWrt software, and the connectivity to the OpenSync™ cloud [8]. OpenSync™ is a cloud-agnostic open-source software that consists of many managers running as separate processes and performing their specific set of tasks. The software code of the Diagnostics Manager (DM) and STATS Manager (SM) was updated for streaming the collected telemetry data and other test results such as the iPerf speed test results via the OpenSync™ cloud. Table 1 summarizes the OpenSync™ manager functionality and their status. Some of the managers listed are required for basic operation, while the other listed managers are optional, depending on the desired functionality.

Table 1: Summary of OpenSync™ Managers' Functionality and Status

Manager Name	Manager Functionality	Manager Status
Diagnostics Manager (DM)	Responsible for spawning the rest of the OpenSync™ managers and optionally monitoring them. It controls starting, stopping, restarting of the OpenSync™ managers, and monitoring the reboot status of the OVSDB. The iPerf speed test software was developed and integrated into the DM such that the speed test can be initiated from the OpenSync™ NOC, and the DS/US speed test results are sent to Grafana dashboard.	Required for basic operation
Connection Manager (CM)	Responsible for establishing the backhaul connection and keeping connectivity to the cloud.	Required for basic operation
Network Manager (NM)	Responsible for managing all network related configuration and network status reporting.	Required for basic system network configuration

Wireless Manager	Not applicable to access CPE devices. Used in Wi-Fi routers to read and updated their configuration and state tables.	Required for basic system network configuration
Queue Manager (QM)	Responsible for aggregating reports from different OpenSync™ Managers	Optional
Statistics Manager (SM)	Responsible for processing all requested wired and wireless statistics and sending results to the cloud. The configuration is done through OVSDb while MQTT is used for the data plane. All the telemetry health metrics mentioned below are collected by write_ctm component as shown in Figure 2, and are transmitted to the SM, which forwards all the collected telemetry data to the QM as shown in Figure 3.	Optional
OpenFlow Manager (OM)	If the OpenVSwitch is used on the device, then the OM is responsible for managing packet flow rules.	Optional
Log Manager (LM)	Responsible for collecting and uploading logs and system information upon the Cloud request (log pull) and for handling log severity setting for running modules.	Optional
Platform Manager (PM)	Responsible for covering specific platform features which can't be covered by other managers such as synchronization between device GUI and cloud, and cloud-managed device parental control.	Optional

4.1. iPerf Speed Test

The iPerf speed test is initiated from the OpenSync™ Network Operations Center (NOC). Submitting a speed test request from the NOC sends a message via Openflow to the access CPE device, and the speed test request is detected on the device by the speed test handler in the OpenSync™ DM. The speed test handler calls a script on the device that in turn invokes an iPerf3 speed test with a pre-defined set of arguments. The speed test is run once to collect the upstream test results, and once again to collect the downstream results. The speed test results from each test are saved to files on the Access CPE device. The STeMTA collectd plugin processes the speed test results from the files and delivers them to the MQTT server, as described in the STeMTA Plugin section. The ability to initiate the iPerf speed test from the OpenSync™ NOC and review the collected speed test results would be very helpful to the Cable operators' call center to quickly address customers' issues.

4.1.1. STeMTA Plugin:

A new speed test plugin, which is called STeMTA, was added to collectd. The STeMTA plugin calls an iPerf speedtest script to initiate the iPerf speed test on WAN port of the D3.1 eMTA. Once an iPerf speed test is completed, then the script writes the downstream and upstream speed test results to iPerf download and upload result files. The STeMTA plugin reads results from these two files and sends them to the write_ctm plugin, which in turn sends the measured speed test results to the OpenSync™ SM.

4.1.2. STLANeMTA Plugin:

A new Speed Test plugin, which is called STLANeMTA, was added to collectd. The STLANeMTA plugin calls iPerf LAN speedtest script to initiate the iPerf speed-test on LAN port of the D3.1 eMTA. Once iPerf speed-test is completed, the script writes the downstream and upstream speed test results to an output file. The STLANeMTA plugin reads results from this file and sends to them to the write_ctm plugin, which in turn sends the measured speed test results to OpenSync™ SM.

4.1.3. Cloud Security

The software architecture also includes several cloud-connectivity security features. The D3.1 eMTA uses OpenSync™ device certificates to authenticate and connect to the OpenSync™ cloud, and to connect to the MQTT server in order to stream the collected telemetry data to the Cable operator's Streaming and Analytics platform. In addition, a shell script in the device is used to monitor all Secure Shell (SSH) and TELNET connections to the D3.1 eMTA, which are reported to the Grafana dashboard.

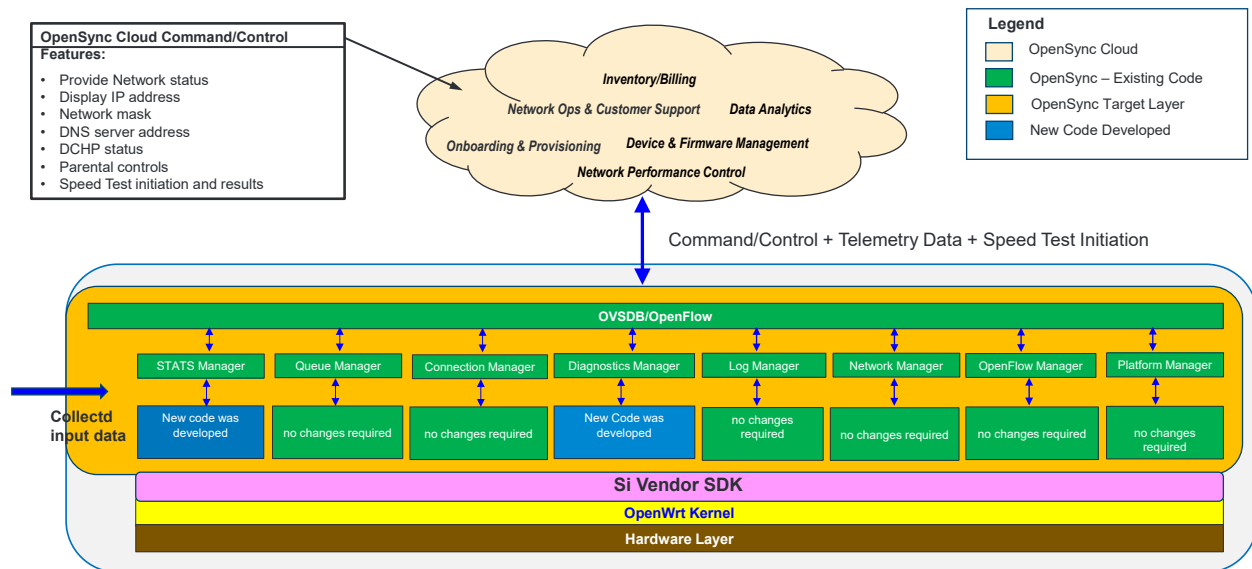


Figure 3: OpenSync™ Software Architecture with the Connectivity to the OpenSync™ Cloud

5. Streaming Telemetry Data Path

Figure 4 shows the telemetry data path (current implementation) from the OpenWrt D3.1 eMTA to the Grafana dashboard via the cable MSO's Streaming and Analytics platform. The collected telemetry data is streamed in Protobuf format to the MQTT server hosted on the OpenSync™ cloud, and then to the cable MSO's Streaming and Analytics platform.

Cable MSO's future network architecture separates the control plane and the data plane as intelligence is no longer resident on hardware devices but rather on the network's software driven controllers where network analytic models can act on traffic behaviors, services flows, and configuration state to predict and

respond in near real-time to the networks changing demands. The network architecture's data plane includes the Cog platform, Data Distribution Bus (DDB), and Data Governance platform. The Cog platform is a data engineering platform that builds enriched data sets called Analytics Data Sets (ADSs), which is represented by the First Normal Form (1NF). These ADSs are distributed across the operator's network, and are used for data modeling and Machine Learning (ML). The DDB, which is shown in Figure 4, is the initial point of data ingestion driven by Apache Kafka [9]. The DDB is also the initial system, where all data is classified as a data asset. The Data Governance platform provides the framework for decisions and accountabilities within the corporate structures to manage and protect the data assets [10]. Any raw data that is not governed as dictated by the data governance standards is transformed for compliance with the standards. As the Cable operators are moving their network's control plane to the cloud, their data plane is maintained in various edge locations deeper into the network, or in this case, the customer's home.

The D3.1 eMTA telemetry data in Protobuf format is ingested by the Kafka Connector source for MQTT, which is part of the data plane of the cable MSO's Streaming and Analytics platform as shown in Figure 4. The Kafka Brokers received the converted telemetry data from Protobuf to Apache Avro™ format, and transmit the telemetry data to the Kafka Connector Sinks [11]. The telemetry data is ingested by different data analytic tools, depending on the type of data. For example, Elasticsearch (ELK) tool ingests time-series data, while MySQL tool ingests relational data before the telemetry data is displayed on the customized Grafana dashboard.

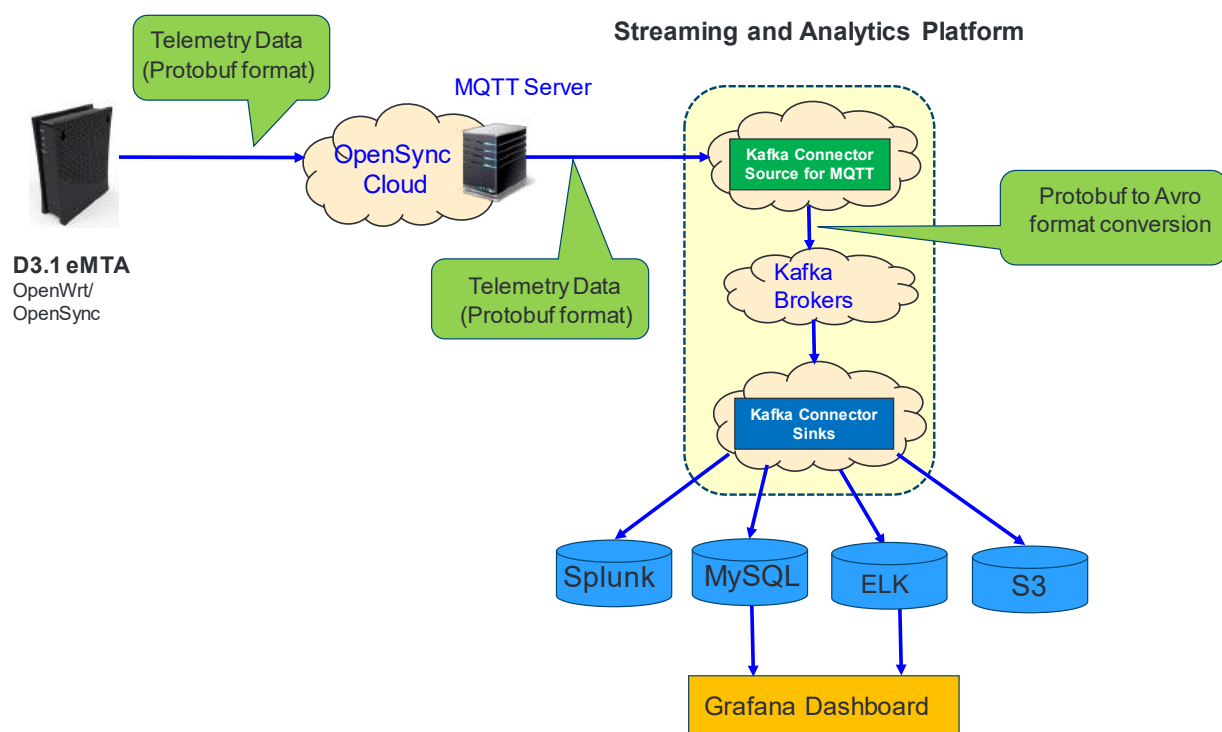


Figure 4: Telemetry Data Path from the Access CPE Device to the Grafana Dashboard

6. Grafana Dashboard Design

The design goal for the customized Grafana dashboard is to concisely present all the different types of telemetry data to the Cable operator's care agent, and to focus the agent attention to any reported failures and unhealthy device metrics. Figure 5 shows, for example, the organized hierarchical color-coded Grafana dashboard with the key components, level 1 health metrics, and their operational status based on pre-determined threshold levels. In addition, the geographical map of the customer location is shown along with reported data telemetry results. The organized Grafana dashboard includes the following information:

- A. D3.1 eMTA Router system information such as:
 - CPU utilization (%) in a given time period
 - Free system memory (%) in a given time period
 - System load (%) in a given time period
 - Networking information such as IP address, network mask, DHCP status, etc.
 - Instantaneous and average system temperature in a given time period
 - Average, minimum, and maximum round-trip IPv4 and IPv6 latency in a selected period of time to the CMTS
- B. Home network traffic from all the wirelessly connected devices via the Pods or Access Points:
 - IPv4/IPv6 of the wirelessly connected client in home network
 - Number of transmitted and received packets for each device
 - IPv4/IPv6 round-trip latency between the D3.1 eMTA and each of the connected clients
- C. Cable modem Downstream/Upstream channel information, including:
 - Downstream channel information (i.e., channel ID, channel type, lock status, channel bonding status, received power level, SNR/MER, channel center frequency, channel width, modulation profile, etc.) – see Figure 9.
 - Upstream channel information (i.e., channel ID, Transmit power level, channel center frequency, channel width, channel bonding status, etc.)
- D. RF downstream and upstream spectrum information – downstream/upstream RF signal power (dBmV) vs. frequency (MHz).
- E. Downstream/Upstream speed test results on the WAN port (i.e., iPerf server in the cable MSO's cloud) and the LAN port (i.e., between iPerf server running on D3.1 eMTA and the connected home network's client).
- F. Security notifications and alarms information, including:
 - Security notifications: for example, if someone is trying to temper with the Access CPE via unauthorized access to the device's management and control GUI via SSH. In this case, the color-coded green status of the security notifications and alarms dashboard component would change to either a color-coded orange or red, indicating an increased security risk.
 - Collected metrics alarms where one or more red thresholds were violated. For example, if the temperature of the device is significantly elevated, and the device is about shut-down or go into energy saving mode.
- G. Customer location map, providing the cable MSO's care agent information where the customer is located within the cable MSO's service area footprint.
- H. Voice health metrics, including:
 - Phone line number
 - phone status (on/off hook)
 - phone line IPv4 and IPv6 addresses

- Voice call start time, end time, call duration, call failed
- phone line registration status
- I. External Battery Backup Unit (EBBU) metrics, including:
 - Manufacturer identity, HW model number, software agent version, battery status
 - EBBU output voltage, and estimated remaining charge capacity,
 - Alarm description (On battery, Low battery, Depleted battery, EBBU shutdown in pending, EBBU shutdown is eminent)
- J. Access CPE device information, including:
 - Cable modem (CM) HW version, SW version, MAC address, serial number
 - IPv4 address and IPv6 address
 - CM system uptime
 - CM security status/type
 - CM connectivity state status/type
- K. Access CPE device reboot information, including:
 - Last device reboot event time, date, and count
 - Last device reboot description
- L. IPv4/IPv6 DOCSIS round-trip latency, including:
 - IPv4 and IPv6 minimum, average, and maximum round-trip latency to the connected CMTS in a given period of time
- M. Device event logs, including any device configuration events, DHCP warnings, DOCSIS events, etc.

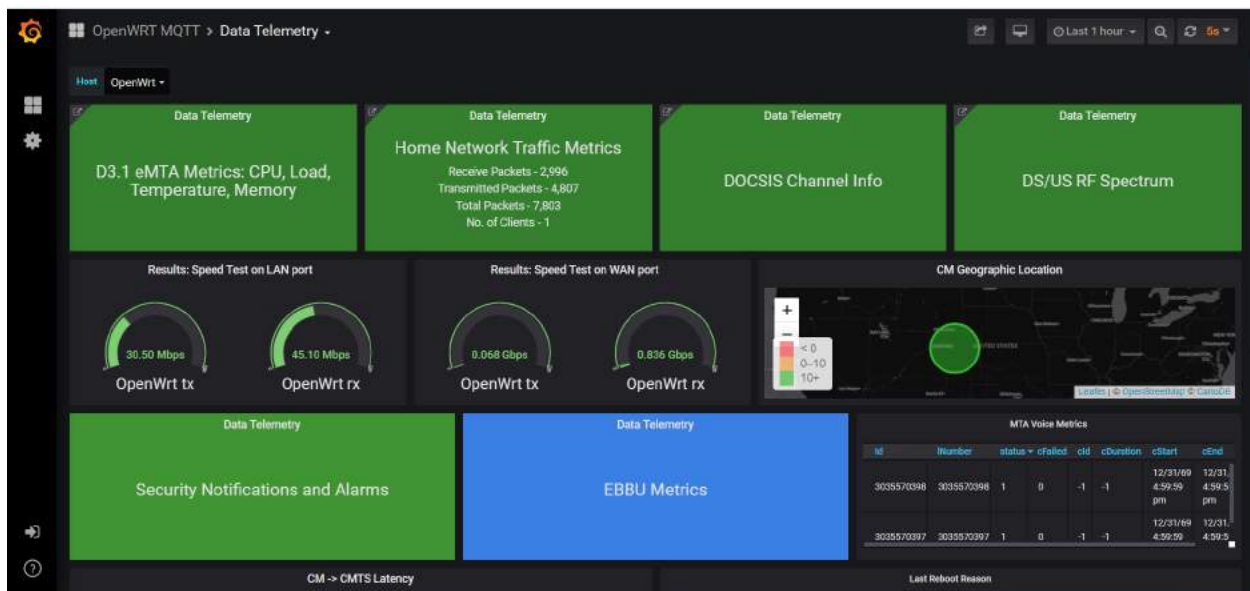


Figure 5: Hierarchical Color-Coded Grafana Dashboard with the Key Telemetry Components

Level 2 telemetry data shows the status of each of the reported health metrics in a given period of time. Figure 6, for example, shows the reported status of the D3.1 system parameter, including system load (%), CPU utilization (%), system temperature (°C), and free memory in the last hour. Each of these reported metrics has a different set of threshold levels to indicate its healthy status. For example, a healthy CPU utilization is below 75%, and it is color-coded green, while unhealthy CPU utilization is above 85%, and it

is color-coded red. CPU utilization between 75% up to 85% is color-coded orange. If all the level 2 D3.1 system parameters are healthy, then the D3.1 eMTA system component (level 1) turns green.




Figure 6: Reported Status Level 2 Access CPE Metrics in the Last Hour

Level 3 telemetry data shows the time behavior of each of the selected metrics in any given time frame. This type of telemetry reporting can be important when deeper insight into the reported metrics is needed to diagnose an issue or abnormal behavior of the D3.1 eMTA device. Figure 6 shows, for example, the reported level 3 telemetry data the time behavior of the CPU utilization reported in the last hour. The CPU utilization data is collected based on a selected time interval, which is 5 seconds in this example. The CPU utilization in this example is low, and varied between about 2.5% to 22.5%. Note that the average system load and CPU utilization are two different things. The system load is a measurement of how many tasks are waiting in a kernel run queue (not just CPU time but also disk activity) over the selected time period. The CPU utilization is a measure of how busy the CPU is during the selected time period.

This type of telemetry data is particularly useful since ML models can be executed on the selected customer data based on the collected historical data to determine if the current reported issue previously occurred, when it occurred, and provide suggested guidelines to the cable MSO's care agent how to mitigate this issue, particularly if this issue previously observed with other customers.




Figure 7: CPU Utilization of Access CPE vs. Time (Level 3) Reported in the Last Hour

One of the challenges for the operator's care agent is to gain visibility into the customer's home network (Figure 1) for network optimization and debugging field issues. This challenge is met by having the access CPE device function as a router, and connected to a Wi-Fi AP via Ethernet cable. Figure 8 shows, for example, the home network traffic parameters (level 2), including the number of transmitted and received packets by each wirelessly connected client in the home network based on their IP address or MAC address reported in the last hour. Instead of using the connected client's IP address or MAC address, the actual client's identification can be displayed on the Grafana dashboard with the integration of a device fingerprinting agent such as a Cujo Artificial Intelligence (AI) agent [12]. The client identification includes device name, device vendor, device model number, device type, device OS, etc. This can be a very useful feature for customers trying to diagnose their home network traffic. For example, customers can identify if there is a specific client that consumes most of the bandwidth in the home network, and/or make changes to their home network configuration.

In addition, the average round-trip latency between the D3.1 eMTA and each of the wirelessly connected client in the last hour is reported. Level 3 home network traffic data can be obtained by selecting a specific client in the home network based on their IP address. For example, the number of transmitted and/or received packets vs. time by the selected client over the selected time interval can be displayed on the Grafana dashboard.

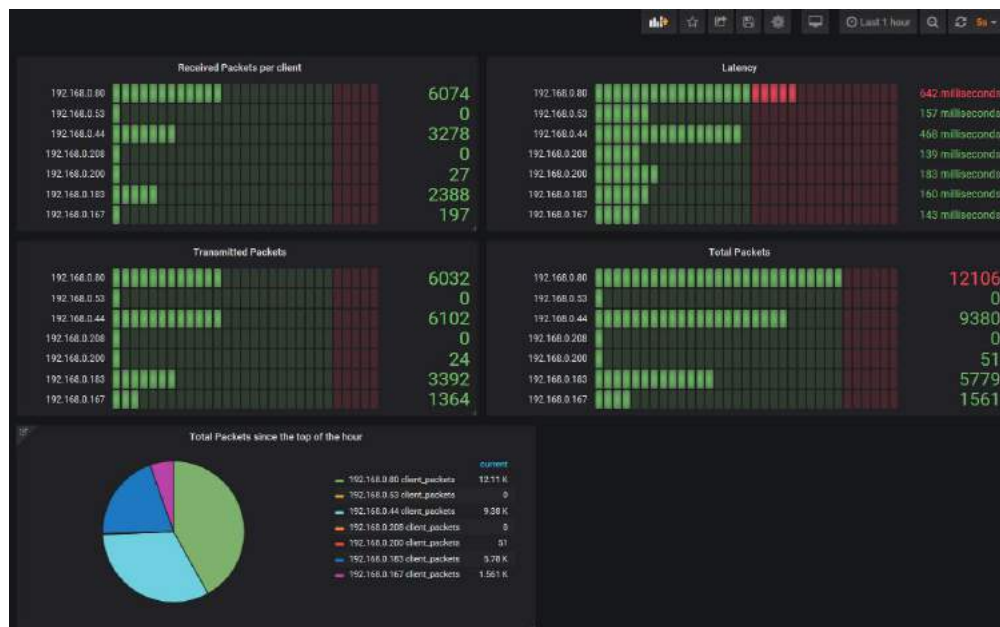


Figure 8: Home Network Traffic Parameters for All the Wirelessly Connected Client in the Home Network

Monitoring the DOCSIS channel information is also important to ensure robust operation of the CM at the customer location. Figure 9 shows, for example, the downstream DOCSIS channel information status, including the channel lock status, channel type, channel bonding status, channel center frequency, channel width, SNR threshold, and received power level, unerrored codewords, number of corrected codewords and uncorrectable codewords. Having access to such monitored historical data with ML models can significantly help to identify troubled CMs in the field compared with other CMs connected to the same CMTS.

OpenWrt Status ▾ System ▾ Network ▾ Logout AUTO REFRESH ON												
Downstream Channel Status												
Channel Index	Channel ID	Lock Status	Channel Type	Bonding Status	Center Frequency	Width	SNR/MER Threshold Value	Receive Level	Modulation/Profile ID	Unerrored Codewords	Corrected Codewords	Uncorrectable Codewords
1	5	Locked	SC-QAM Downstream	Bonded	627000000Hz	6000000Hz	47.2dB	5.9dBmV		18446744072098527978	0	0
2	1	Locked	SC-QAM Downstream	Bonded	603000000Hz	6000000Hz	47.6dB	6.2dBmV		18446744072096025101	0	0
3	2	Locked	SC-QAM Downstream	Bonded	609000000Hz	6000000Hz	47.6dB	6.2dBmV		18446744072096020855	0	0
4	3	Locked	SC-QAM Downstream	Bonded	615000000Hz	6000000Hz	47.2dB	6.1dBmV		18446744072096045564	0	0
5	4	Locked	SC-QAM Downstream	Bonded	621000000Hz	6000000Hz	47.4dB	6dBmV		18446744072096039927	0	0
6	6	Locked	SC-QAM Downstream	Bonded	633000000Hz	6000000Hz	47.1dB	5.9dBmV		18446744072096061356	0	0
7	7	Locked	SC-QAM Downstream	Bonded	639000000Hz	6000000Hz	47.1dB	5.9dBmV		18446744072096074839	0	0
8	8	Locked	SC-QAM Downstream	Bonded	645000000Hz	6000000Hz	47.1dB	5.8dBmV		18446744072096093890	0	0
9	9	Locked	SC-QAM Downstream	Bonded	651000000Hz	6000000Hz	46.9dB	5.8dBmV		18446744072096107317	0	0

Figure 9: Downstream DOCSIS Channel Information Status

7. Comparison with Other Streaming Telemetry Methods

There are other streaming telemetry data methods for access CPE devices. One of the streaming methods is based on the Internet Protocol Detail Record Streaming Protocol (IPDR/SP). The Cable Modem Termination System (CMTS) is using the IPDR/SP via CableLabs-defined schemas to collect customer data usage via billing records from specific service flows used by the CMs and export them to IPDR Collectors [13]. IPDR/SP utilizes the concept of templates in order to eliminate the transmission of redundant information such as field identifiers and typing information on a per data record basis.

Specifically, IPDR/SP Subscriber Account Management Interface Specification (SAMIS) Type 1 schema is probably the most common IPDR schema in use by the cable MSOs. SAMIS Type 1 uniquely identifies the specific CM attributes and service flow's attributes serviced by the CM. Currently, these billing records are collected every 15 minutes by the CMTS, which forward the collected records to the IPDR Collectors hosted in the regional data centers before the data is ingested by the cable MSO's Streaming and Analytics platform. Expanding the CMTS usage of the IPDR/SP to collect all of the various performance and health metrics, alarms, and notifications from each field-deployed CM would overburden each CMTS with huge amounts of data. To estimate the magnitude of this issue, it is assumed that each CM streams ≈ 62.5 MB telemetry data every 24 hours based on the current implementation, and each CMTS is connected to ≈ 20 k CMs. Consequently, each CMTS would receive about 1.25TB of telemetry data every 24 hours. Thus, this approach is burdensome since the CMTS do not process or make decisions on the enormous amount of collected telemetry data. Furthermore, new IPDR schemas for non-DOCSIS parameters such as voice metrics would need to be defined, standardized, and integrated with the access CPE firmware.

Model-Driven Telemetry (MDT) is another modern method for continuously streaming operational data from network devices such as the access CPE using a push model. Applications need to subscribe to a set of the access CPE device's Yet Another Next Generation (YANG) data models over standard protocols,

and push the collected telemetry data from the device when a change has occurred. MDT is not new to the cable industry as several Cable operators have already implemented and deployed MDT data collection and monitoring systems in their network [14]. Implementing MDT on access CPE devices requires the development a new software layer and components for YANG data models. In addition, a common set of standard Application Programming Interfaces (APIs) need to be defined for integration with an OpenWrt-based access CPE software stack. Furthermore, no detailed telemetry comparison analysis of performance vs. cost has been done to justify such an MDT-based development effort by the Cable operators.

In contrast, the OpenWrt and OpenSync™-based streaming telemetry method uses an agile lightweight and efficient smart agent based on an open-source code with customized plugins using the Silicon vendor's APIs. The proposed approach is to segregate the telemetry data such that all customer's billing records continue to be provided using the IPDR/SP, while all of the other access CPE device's performance and health metrics are directly transmitted to the cable MSO's Streaming and Analytics platform via the OpenSync™ cloud. Consequently, the CMTSs are not overburdened with huge amounts telemetry data.

8. Conclusion

In this paper, an agile OpenWrt software stack integrated with OpenSync™ layer and Silicon vendor SDK with carrier-grade IPv4/IPv6 routing functionality was developed on a common existing access CPE hardware. The connectivity of this access CPE device to the OpenSync™ cloud provides a standardized command and control method for networking services as well as operator-friendly services such as onboarding and provisioning on field-deployed devices, device firmware management, network operations and customer support, billing and inventory support. A smart remote agent was developed and integrated with the OpenWrt software stack that enables the access CPE device to stream various types of telemetry data to the cable MSO's Streaming and Analytics platform via the MQTT server hosted on the OpenSync™ cloud for analysis, and displays the collected data on a hierarchical color-coded Grafana dashboard. The streaming telemetry data consists of a wide variety of information, including:

- Access CPE system information
- CM device information
- Home network traffic information from all the wirelessly connected clients
- DS/US DOCSIS channel information
- DS/US RF spectrum output
- Event and alarms information for the collected metrics
- Speed test results on both the WAN and LAN ports
- Voice metrics information
- EBBU status information

Comparison with other streaming telemetry methods such as IPDR/SP and MDT reveal various challenges with the implementation of these methods. For example, expanding the IPDR/SP usage by the CMTS to collect all of the various performance and health metrics, alarms, and notifications would overburden each CMTS with huge amounts of data that it does not process or make decisions on is not an attractive approach. Implementing MDT on access CPE devices requires the development a new software layer and components for YANG data models without clear benefits.

Finally, as shown in this paper, the adoption of an OpenWrt-based streaming telemetry offers clear benefits to Cable operators. First, it enables cloud-based management of the access CPE devices and direct streaming of the telemetry data to the cable MSO's Streaming and Analytics platform via the OpenSync™ cloud without overburdening the CMTS. Second, the availability of the telemetry data to the operator's

care agent is likely to enhance customer satisfaction by reducing the number of truck rolls as field issues are resolved more quickly. Third, collaborations among different Cable operators will enable the industry to standardize the OpenWrt-based software architecture with the streaming telemetry across different types of CPE hardware platforms. Furthermore, the standardized agile software stack is expected to accelerate the development and deployment of new revenue-generating services.

Acknowledgment

The authors would like to acknowledge technical support from Jay Liew and Philip Anderson and management support from Ahmad Ansari and Matt Petersen at Charter Communications.

Abbreviations

Table 2: Abbreviations Table

Acronym	Stand For
AP	Access Point
API	Application Programming Interface
CM	Cable Modem
CPE	Customer Premise Equipment
CPU	Central Processing Unit
CMTS	Cable Modem Termination System
DDB	Data Distribution Bus
DOCSIS	Data over Cable System Interface Specification
EBBU	External Battery Backup Unit
eMTA	Embedded Media Terminal Adapter
FTTH	Fiber To The Home
ICMP	Internet Control Message Protocol
IPDR	Internet Protocol Detail Record
LAN	Local Area Network
MDT	Model Driven Telemetry
MER	Modulation Error Ratio
ML	Machine Learning
MQTT	Message Queue Telemetry Transport
MSO	Multiple System Operator
NOC	Network Operations Center
OEM	Original Equipment Manufacturer
ONU	Optical Network Unit
OS	Operating System
OVSDB	Open vSwitch Data Base
QM	Queue Manager
RF	Radio Frequency
SAMIS	Subscriber Account Management Interface Specification
SDK	Software Development Kit
SM	STATS Manager
SNR	Signal to Noise Ratio
WAN	Wide Area Network

Bibliography & References

- [1] U.S. Smart Homes, Statista.
<https://www.statista.com/outlook/279/109/smart-home/united-states>
- [2] <https://openwrt.org/>
- [3] ISO/IEC 20922:2016, Information Technology – Message Queuing Telemetry Transport (MQTT) v3.1.1, <https://www.iso.org/standard/69466.html>.
- [4] <https://grafana.com/docs/grafana/latest/features/dashboard/dashboards/>
- [5] <https://collectd.org/>
- [6] https://openwrt.org/docs/guide-user/services/network_monitoring/wrtbwmon
- [7] <https://www.opensync.io/>
- [8] <https://svn.nmap.org/nmap/COPYING>
- [9] <https://kafka.apache.org/>
- [10] K.Wende, A Model for Data Governance – Organising Accountabilities for Data Quality Management, Association for Information Systems(2007).
<https://aisel.aisnet.org/cgi/viewcontent.cgi?article=1079&context=acis2007>
- [11] <https://avro.apache.org/docs/current/>
- [12] <https://cujo.com/agent/>
- [13] Data-Over-Cable Service Interface Specifications (DOCSIS) 3.1, CCAP Operations Support System Interface Specification, CM-SP-CCAP-OSSIV3.1-I16-190917.
- [14] P. Sowinski, A. Smith, and T. Liu, Remote PHY 2.0, the Next Steps for Remote PHY Technology, SCTE Technical Papers (2019).

DOCSIS 4.0 Network Migration Made Easy

A Technical Paper prepared for SCTE•ISBE by

Ayham Al-Banna, Ph.D.

Director of Product Line Management & Fellow
CommScope
2400 Ogden Ave., Suite 180, Lisle, IL 60532, USA
630-281-3009
Ayham.Al-Banna@CommScope.com

Tom Cloonan, Ph.D.

CTO – Network Solutions
CommScope
2400 Ogden Ave., Suite 180, Lisle, IL 60532, USA
630-281-3050
Tom.Cloonan@CommScope.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Latest Traffic Engineering Trends & Potential COVID-19 Effect.....	4
3. Network Migration Tools.....	6
3.1. Selective Subscriber Migration.....	6
3.2. Node Splits and Node Segmentations.....	6
3.3. Digital Video, Switched Digital Video & IPTV.....	7
3.4. Split Upgrade.....	7
3.5. Full Duplex DOCSIS.....	7
3.6. Dynamic Soft-FDD.....	7
3.7. Extended Spectrum DOCSIS.....	8
3.8. Active Taps.....	8
3.9. Fiber To The Tap.....	8
4. Interactions of Traditional FDD, FDX & Dynamic Soft-FDD, and ESD.....	8
5. Time-Aware Decision Making.....	12
6. Example Network Migration Strategy.....	20
7. Conclusions.....	22
Abbreviations.....	22
Bibliography & References.....	23

List of Figures

Title	Page Number
Figure 1 - DS Tavgl latest statistics before COVID-19.....	5
Figure 2 - US Tavgl latest statistics before COVID-19.....	5
Figure 3 - Decision tree combining multiple interrelated network migration tools: traditional FDD with DOCSIS 3.0 splits, DOCSIS 3.1/DOCSIS4.0 split change, Dynamic Soft-FDD, and ESD.....	9
Figure 4 - US & DS Tmax values that can be offered over time for various split options and DS spectrum sizes. (SG size = 250, 66 Digital video channels).....	14
Figure 5 - US & DS Tmax values that can be offered over time for various split options and DS spectrum sizes. (SG size = 125, 66 Digital video channels).....	15
Figure 6 - US & DS Tmax values that can be offered over time for various split options and DS spectrum sizes. (SG size = 64, 66 Digital video channels).....	16
Figure 7 - US & DS Tmax values that can be offered over time for various split options and DS spectrum sizes. (SG size = 250, IPTV video channels).....	17
Figure 8 - US & DS Tmax values that can be offered over time for various split options and DS spectrum sizes. (SG size = 125, IPTV video channels).....	18
Figure 9 - US & DS Tmax values that can be offered over time for various split options and DS spectrum sizes. (SG size = 64, IPTV video channels).....	19

List of Tables

Title	Page Number
Table 1 - US & DS Tmax values that can be supported with today's networks (different configurations)	11
Table 2 - US & DS Tmax values that can be supported via different migration paths (with Digital video)	11

1. Introduction

The plethora of new technologies and alternatives has made the task of HFC network migration more challenging than ever! The MSOs are faced with difficult decisions as their strategy executives and HFC network architects try to navigate the future of their HFC networks. The difficulty comes from identifying potential market challenges and addressing those challenges using the right technology in a timely and cost-effective manner.

DOCSIS 4.0 offers various technologies that will enable the delivery of 10 Gbps DS peak rates and 5 Gbps US peak rates. Utilizing the right technology at the right time will be instrumental for the success of the MSOs. This paper attempts to provide the list of available tools in the network migration toolkit and proposes a time-aware methodology for using these tools to yield a cost-effective migration strategy that meets customers traffic demand and addresses competition.

This paper is organized as follows: Section 2 introduces the latest traffic engineering trends and potential COVID-19 impact. Various network migration tools are described in Section 3. Section 4 discusses the decision interactions between some of the migration tools like traditional FDD, FDX & dynamic soft-FDD, and ESD. Time-aware decision process is discussed in Section 5. Section 6 provides an example migration strategy that uses various migration tools. Finally, the paper is concluded in Section 7.

2. Latest Traffic Engineering Trends & Potential COVID-19 Effect

Figure 1 and Figure 2 show the latest busy-hour average subscriber consumption rate (Tavg) for the DS & US directions, respectively, collected from multiple MSOs. These trends, which were captured in the beginning of 2020 and before the COVID-19 pandemic effect has spread, show that Tavg averaged across various MSOs yielded 2.36 Mbps for the DS and 164 kbps for the US. Additionally, the 2020 statistics show that the 3-year DS Tavg CAGR had dropped from 34% to 30% and the US Tavg CAGR had slightly dropped from 22% to 21%.




Figure 1 - DS Tavg latest statistics before COVID-19




Figure 2 - US Tavg latest statistics before COVID-19

At the outset of the COVID-19 pandemic, the DS and US consumption increased significantly as countries, states, and cities enforced a lockdown to control the disease. With many people at home and the

lack of outside activities, increased network consumption was in the form of video streaming, working-from-home, virtual video gatherings, gaming, etc. The increase in the US utilization was about 35% and the increase in the DS utilization was 20%. As some countries, states, and cities started to reduce the limits of the lockdown, the consumption (relative to pre-COVID numbers) has decreased recently (as of July 4, 2020) and currently shows an increase of about 22% for the US and 14% for the DS [NCTA-COVID-19-Tracker-2020]. The US & DS will not likely go back to the original levels because of the “new normal” after COVID-19, where many people will likely continue to work from home. It is probably reasonable to assume that the US will settle to 18% (half of 35%) and the DS will settle to 10% (half of 20%). Therefore, our traffic engineering assumptions for the analyses provided later in this paper assume an US Tav_g of 190 kbps (instead of 164 kbps) for 2020 and a DS Tav_g of 2.6 Mbps (instead of 2.36 Mbps) for 2020. Note that these percentages represent one-time jump in Tav_g due to COVID-19 and are independent from the CAGR percentages discussed in the beginning of this section, which represent the annual growth in Tav_g.

3. Network Migration Tools

This section briefly describes the available tools in the network migration toolbox. These tools can be used to gradually and concurrently migrate along an optimal migration path that extends the useful life-span of today’s HFC network while gradually transforming it into the desired end goal as a FTTH network.

3.1. Selective Subscriber Migration

Selective subscriber migration refers to the concept of performing HFC surgical operations, whereby customers demanding very high peak rates are moved to a different platform, such as an overlay FTTH network running a specific flavor of PON. This concept, which was explained in detail in [Spring-Forum-Migration-2016], avoids the need to upgrade the whole network in order to meet the abnormal service levels associated with very small penetration rates. For instance, if only a select number of residential or business customers request symmetrical 10 Gbps service, while the rest of customers are happy with sub 1 Gbps service, then it makes sense to move those few customers to 10G PON (e.g., XGS-PON or 10G-EPON) as opposed to overhauling the whole HFC network. This tool can be used to constrain the goal of HFC network migration to address rates offered by EPON or GPON but not with 10G PON competition. Using this constraint in the decision process will yield an optimal cost-effective competitive migration strategy.

3.2. Node Splits and Node Segmentations

Node splits and node segmentations provide another powerful tool within the HFC network migration toolkit. They can help reduce the number of subscribers per SG, which will lead to less congestion and therefore the ability to support higher peak service rates for a longer time. Recall that the total capacity required by a SG that has N_{sub} subscribers, each consuming an average busy hour rate of Tav_g, and requesting a peak rate service of T_{max}, is given by the following formula, where K is a QoE coefficient that is recommended to be between 1-1.2 to provide good QoE:

$$\text{Required SG Capacity} = N_{\text{sub}} \cdot \text{Tavg} + K \cdot \text{Tmax}$$

Note that with fixed available capacity, reducing the number of subscribers using a node split will reduce the first term of the above equation, which will enable supporting higher T_{max} (within the second term) using the fixed available capacity.

3.3. Digital Video, Switched Digital Video & IPTV

The use of analog video channels comes at a price because they do not use the spectrum very inefficiently. Moving video from analog to digital channels will enable better utilization of the scarce spectrum via more efficient encoding schemes of the digital content. Adding SDV capabilities that only transmit video streams when being viewed can also help reduce spectrum utilization for video services. As time goes on, reducing the number of digital video channels and migrating to IPTV video (which uses more spectrally-efficient DOCSIS 3.1 channels) will provide yet another level of efficient usage of spectrum.

3.4. Split Upgrade

There are multiple US split options supported by the DOCSIS3.1 specifications: sub-split (5-42 MHz in NA and 5-65 MHz in Europe), Mid-split (5-85 MHz), and High-split (5-204 MHz). Supporting US peak rates up to 400 Mbps will require moving from sub-split to mid-split. On the other hand, supporting US peak rates in excess of 1 Gbps will require a move to a high-split architecture.

DOCSIS 4.0 specifications add other ultra-split options with US limits up to 300 MHz, 396 MHz, 492 MHz, and 684 MHz. These splits can be used in a fixed or dynamic manner as will be explained later in this paper.

3.5. Full Duplex DOCSIS

FDX is a technology that is designed to allow the US & DS traffic to share the same spectrum simultaneously. The FDX technology was thoroughly explained in [SCTE-Tec-Soft-FDD-2019]. It is optimized for DAA N+0 network architecture. MSOs who find it expensive to migrate their networks to N+0 will likely need other alternatives. One of those alternatives is to use FDX amplifiers. However, those amplifiers will cause the problem of Interference Group elongation [SCTE-Tec-Soft-FDD-2019], where most of the subscribers become a member of the same interference group (i.e., interfering with each other); therefore true FDX operation in N+x networks may not be feasible.

3.6. Dynamic Soft-FDD

In the previous section, the challenge of running FDX in N+x ($x > 0$) networks was described. A potential solution for that problem is to run the system in Dynamic Soft-FDD mode [SCTE-Tec-Soft-FDD-2019], where the split is changed dynamically to match the traffic demand and hence offering the same benefits as FDX, but with typically larger Interference Groups. In a nutshell, Dynamic Soft-FDD is viewed as the tool that enables FDX operation in cascaded networks. Dynamic Soft-FDD is based on the same modem silicon as FDX and can help avoid replacing taps and passives beyond 1.2 GHz. However, more complex amplifiers will be needed, which may put a limit on the maximum cascade depth that can be supported with Dynamic Soft-FDD. Additionally, operating the Dynamic Soft-FDD plant could be challenging for some MSOs.

3.7. Extended Spectrum DOCSIS

Another tool in the network migration toolbox is ESD, where the DS spectrum is allowed to go beyond 1.2 GHz. DOCSIS 4.0 introduced requirements for 1.8 GHz operation in equipment whose housing supports 3 GHz spectrum. The ESD technology was described in detail in [SCTE-Tec-Soft-FDD-2019] [Spring-Forum-ESD-2016]. While the ESD technology requires changing taps and amplifiers, it works in an FDD mode just like today's networks. The ESD amplifier design can also be challenging due to the high gain requirements needed for operation at high frequencies, where attenuation is significant. Additionally, it should be noted that the limited TCP can be a potential issue with very long plants but not so much with short or medium plants, which make the majority of current HFC networks. Those challenges and potential solutions were discussed in [SCTE-Tec-Soft-FDD-2019].

3.8. Active Taps

As the networks get deeper in fiber deployment, they may eventually get to N+0. Before pulling fiber any deeper, the concept of active taps could be beneficial as taps are transformed into small active devices that support relatively small gain values. This enables continued use of the hardlines with higher modulation orders at higher frequencies, which will yield additional capacities [ANGA-Cable-Migration-2019] [SCTE-Tec-Soft-FDD-2019].

3.9. Fiber To The Tap

FTTT is a natural late-stage step to take before pulling fiber to the home, whether active taps were used or not as an intermediate step. FTTT enables the use of existing drop cables that can support frequencies up to 25+ GHz, which translates to data rates of more than 200 Gbps [Spring-Forum-ESD-2016]. In fact, when FTTT is used, the coaxial cable network becomes a point-to-point network which will enable ESD operation to be used in combination with FDX operation (without undergoing the current complexities of FDX that are faced with the current multi-point-to-single-point HFC networks).

4. Interactions of Traditional FDD, FDX & Dynamic Soft-FDD, and ESD

It can be confusing to think of multiple interrelated tools at the same time. In order to address this challenge, a decision tree was developed for N+x networks as shown in Figure 3. Observe that the decision tree considers multiple interrelated network migration tools including traditional FDD with DOCSIS 3.0 splits, DOCSIS 3.1/DOCSIS4.0 split change, Dynamic Soft-FDD, and ESD. In this context, recall that Dynamic Soft-FDD is the FDX flavor for N+x networks.

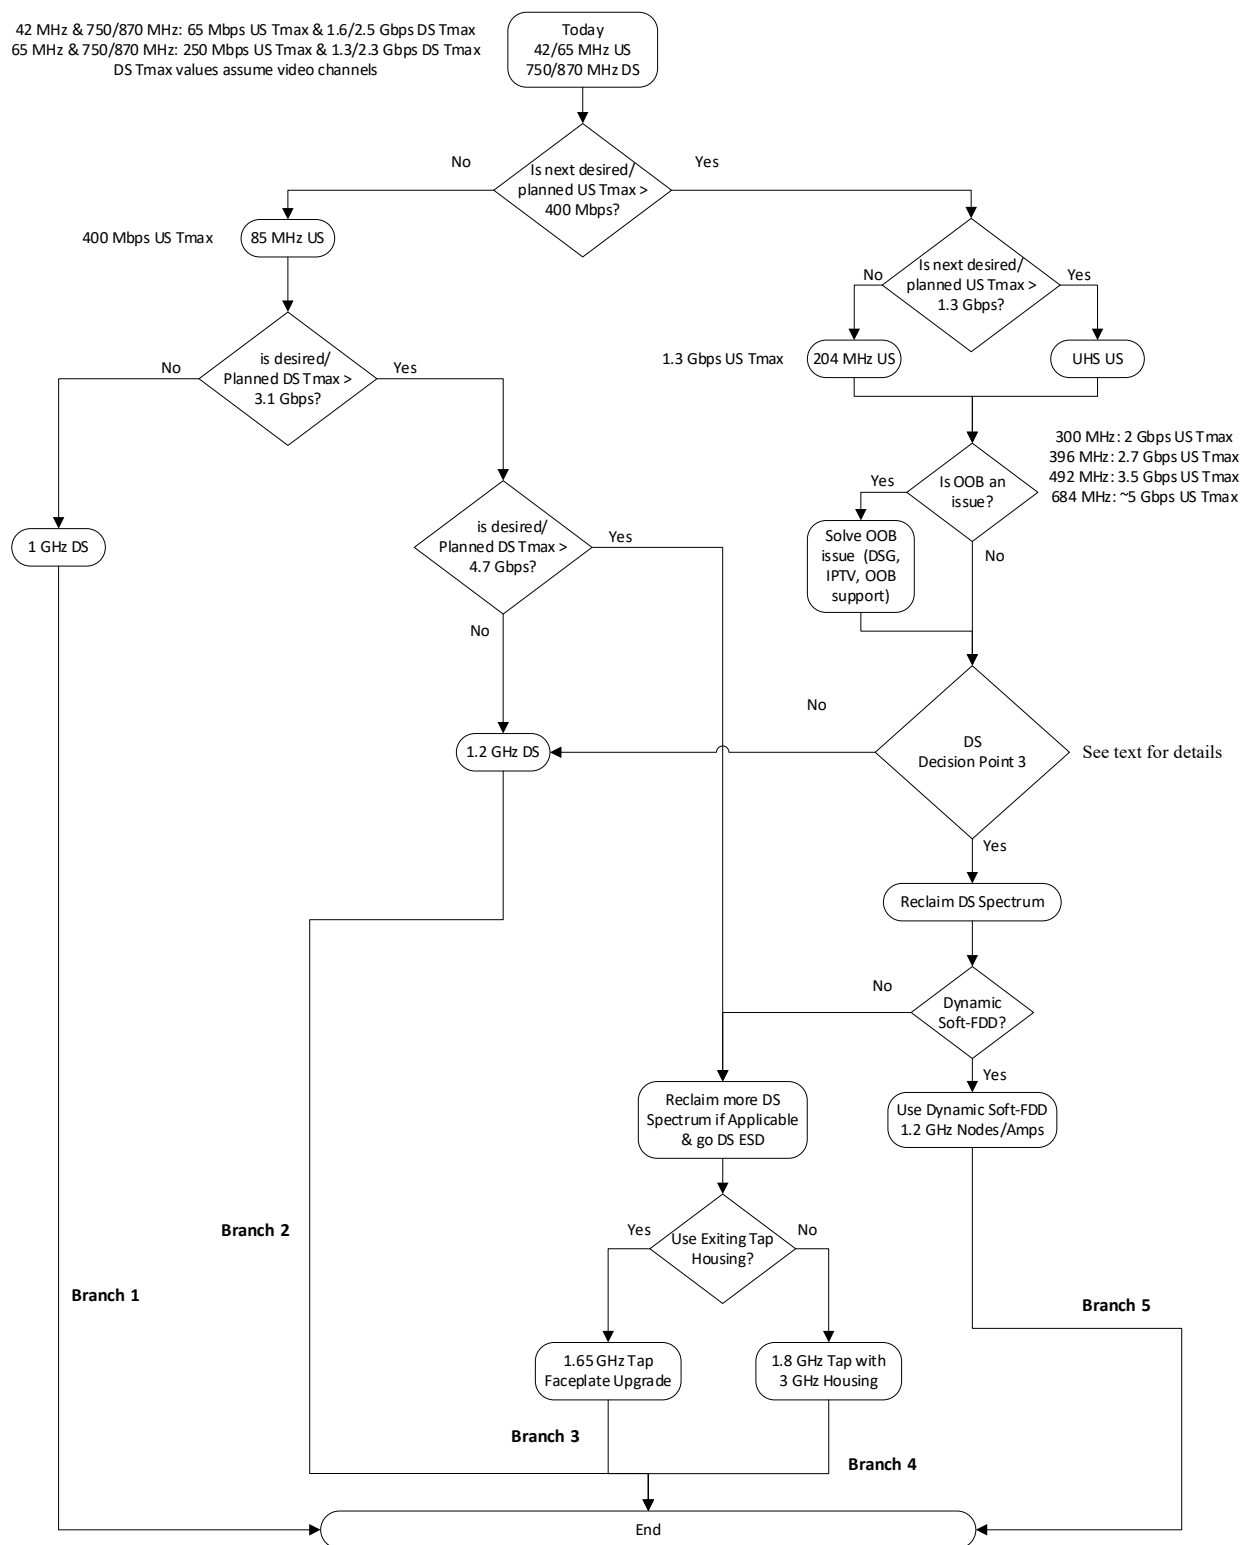


Figure 3 - Decision tree combining multiple interrelated network migration tools: traditional FDD with DOCSIS 3.0 splits, DOCSIS 3.1/DOCSIS4.0 split change, Dynamic Soft-FDD, and ESD

The ‘DS Decision Point 3’ in Figure 3 contains a list of decision questions, depending on the chosen US split, as follows:

Case (Split)

- Case A: 204 MHz Split
is desired/Planned DS T_{max} > 3.6 Gbps?
- Case B: 300 MHz Split
is desired/Planned DS T_{max} > 2.7 Gbps?
- Case C: 396 MHz Split
is desired/Planned DS T_{max} > 1.7 Gbps?
- Case D: 492 MHz Split
is desired/Planned DS T_{max} > 900 Mbps?
- Case E: 684 MHz Split
is desired/Planned DS T_{max} > 0 Gbps?

It can be observed from the above decision point case statement that 1.2 GHz systems cannot support DS T_{max} values more than 900 Mbps if the US split is 492 MHz or higher. Note that the 1.65 GHz Tap faceplate upgrade option refers to reusing existing taps housings by only replacing the faceplate with a new faceplate that supports higher frequencies, which can help in reducing the labor costs.

The assumptions that were used to calculate the T_{max} numbers in the decision tree are as follows:

- **Traffic Engineering**
 - K = 1
 - N_{sub} = 250
 - US T_{avg} = 0.19 Mbps (0.164 Mbps with added step increase of 18% due to COVID-19)
 - DS T_{avg} = 2.6 Mbps (2.36 Mbps with added step increase of 10% due to COVID-19)
- **US**
 - # of 6.4 MHz SC-QAM channels = 4
 - SC-QAM spectral efficiency = 4.15 bps/Hz
 - OFDMA spectral efficiency = 7.5 bps/Hz
 - US spectrum start frequency = 15 MHz
- **DS**
 - # of 6 MHz SC-QAM channels = 20
 - SC-QAM spectral efficiency = 6.33 bps/Hz
 - OFDM spectral efficiency = 7.8 bps/Hz
 - Video: 66 Digital channels

Clearly, it should be noted that similar decision trees can easily be created for a different set of assumptions.

Note that analyses used in developing the decision tree assumed a SG size of 250 subscribers and 66 digital video channels, which can be argued to be good representative values of today’s HFC networks. Table 1 provides the T_{max} values that can be supported over today’s HFC networks assuming no modifications. On the other hand, Table 2 provides the T_{max} values that can be supported over a modified network using the different migration branches that were illustrated in Figure 3, which are:

1. Branch 1: 85 MHz US with 1 GHz DS
2. Branch 2: 1.2 GHz DS (with different US splits)
3. Branch 3: 1.6 GHz DS using tap faceplate upgrade (~ESD with different US splits)
4. Branch 4: 1.8 GHz DS using new ESD taps (ESD with different US splits)

5. Branch 5: 1.2 GHz Dynamic Soft-FDD

Table 1 - US & DS Tmax values that can be supported with today's networks (different configurations)

US Split (MHz)	Top of DS Spectrum (MHz)	US Tmax (Mbps)	DS Tmax (Gbps)
42	750	60	1.5
42	870	60	2.5
42	1002	60	3.5
65	750	250	1.3
65	870	250	2.2
65	1002	250	3.2

Table 2 - US & DS Tmax values that can be supported via different migration paths (with Digital video)

US Split (MHz)	US Tmax (Gbps)	DS Start Frequency (MHz)	Branch 1 FDD 1.0 GHz DS Tmax (Gbps)	Branch 2 FDD 1.2 GHz DS Tmax (Gbps)	Branch 3 FDD 1.65 GHz DS Tmax (Gbps)	Branch 4 FDD 1.8 GHz DS Tmax (Gbps)	Branch 5 Dyn. Soft-FDD 1.2 GHz DS Tmax (Gbps)
85	0.4	108	3.1	4.7	NA	NA	4.7
204	1.3	258	NA	3.6	6.9	8.1	4.7
300	2.0	372	NA	2.7	6.1	7.2	4.7
396	2.7	492	NA	1.7	5.1	6.2	4.7
492	3.5	606	NA	0.9	4.2	5.4	4.7
684	5.0	834	NA	0	2.5	3.6	4.7

It should be noted that the above decision tree must be used along with other tools like selective subscriber migration, node splits, migrating to IPTV, etc. This wholistic approach to the decision-making process yields an optimal migration path from complexity and cost point views. In particular, the table above contains technology options that are either costly or not available today, which should be taken into consideration as the MSOs plan a stepping-stone migration process.

5. Time-Aware Decision Making

A key missing aspect from the decision tree and the other tools described in the previous section is the time dimension of the decision. Specifically, the busy hour per subscriber consumption rate (i.e., T_{avg}) continues to grow as time goes on, which in turn reduces the T_{max} values that can be supported given a particular network architecture. Also, the SG size (N_{sub}) decreases as more node splits are undertaken, which leads to higher T_{max} values. However, these two terms are multiplied by each other in the QoE formula so the net effect may not be obvious unless simulations are performed. Another aspect that changes with time is the migration of digital video channels to IPTV which leads to more efficient video delivery and that yields support for higher T_{max} values. This section introduces the time-awareness aspect into the decision process.

In order to take the above time-affected tools into consideration, the analysis in this section uses the latest traffic CAGR numbers presented in section 2 (i.e., US CAGR of 21% and DS CAGR of 30%). Also, the impact of node splits over time is studied by analyzing the effect of moving to smaller SG sizes (i.e., reducing the SG size from 250 to 125 and then to 64). Finally, moving video delivery from 66 digital channels to IPTV is also simulated to understand the effect of this move on the life of the network. Finally, multiple curves are illustrated to show the impact of changing the US split and moving the DS top end to 1.65 GHz or 1.8 GHz.

The Time-aware simulation assumptions are as follows:

- **Traffic Engineering**
 - $K = 1$
 - $N_{sub} = 250$ subscribers (Penetration ratio of data subscribers 50%)
 - US $T_{avg} = 0.19$ Mbps (0.164 Mbps with added step increase of 18% due to COVID-19) in 2020
 - US CAGR = 21%
 - DS $T_{avg} = 2.6$ Mbps (2.36 Mbps with added step increase of 10% due to COVID-19) in 2020
 - DS CAGR = 30%
- **US**
 - # of 6.4 MHz SC-QAM channels = 4
 - SC-QAM spectral efficiency = 4.15 bps/Hz
 - OFDMA spectral efficiency = 7.5 bps/Hz
 - US spectrum start frequency = 15 MHz
- **DS**
 - # of 6 MHz SC-QAM channels = 20
 - SC-QAM spectral efficiency = 6.33 bps/Hz
 - OFDM spectral efficiency = 7.8 bps/Hz
 - Video options: 66 Digital channels, or IPTV
- **IPTV Video**
 - Penetration ratio of video subscribers = 30%
 - Unicast only
 - HD MPEG4 bit rate = 5 Mbps
 - UHD/4K bit rate = 17 Mbps

- 90% HD / 10% UHD mix
- 5% VOD

The results of the analyses are shown in Figure 4 - Figure 9. The first three figures assume digital video channels and study the effect of node splits by varying the SG size from 250, to 125, and finally to 64. Similarly, the last three figures change the SG size but assuming IPTV video delivery instead of digital video channels.

Understanding how to interpret the curves can be best illustrated using an example. For instance, in Figure 4, curves with no markers show the highest peak service rate (Tmax) that can be supported in the US direction for different US split options. On the other hand, curves with markers indicate the highest DS Tmax that can be supported for different configurations (particular US split combined with DS spectrum limit).

Let's assume that the MSO would like to offer symmetrical 1 Gbps service, it can be seen from the Figure 4 that high-split (i.e., 204 MHz US split) can support US Tmax of 1 Gbps until about 2031 – this can be observed from the non-marker gray curve that corresponds to 'US Tmax: 204 MHz US Split'. Using the same approach, it can be seen that a 204 MHz US with 1.2 GHz DS system can offer a DS Tmax value of 1 Gbps beyond 2026 – this can be observed from the circle-marker dark blue curve that corresponds to 'DS Tmax: 204 MHz Split with 1.2 GHz DS'.

The effect of a node split that divides the SG size in half (from 250 to 125), can be observed in Figure 5, where it can be seen that high-split can offer US Tmax of 1 Gbps to about 2034 and high-split 1.2 GHz DS system can offer a DS Tmax value of 1 Gbps about 2029. Also, observe how a move to IPTV can further extend the lifespan of the network. For example, for a SG of 125, a move to IPTV will enable a high-split network with 1.2 GHz DS to support DS Tmax value of 1 Gbps to about 2031 as can be observed in Figure 8.




Figure 4 - US & DS Tmax values that can be offered over time for various split options and DS spectrum sizes. (SG size = 250, 66 Digital video channels)




Figure 5 - US & DS Tmax values that can be offered over time for various split options and DS spectrum sizes. (SG size = 125, 66 Digital video channels)




Figure 6 - US & DS Tmax values that can be offered over time for various split options and DS spectrum sizes. (SG size = 64, 66 Digital video channels)




Figure 7 - US & DS Tmax values that can be offered over time for various split options and DS spectrum sizes. (SG size = 250, IPTV video channels)




Figure 8 - US & DS Tmax values that can be offered over time for various split options and DS spectrum sizes. (SG size = 125, IPTV video channels)




Figure 9 - US & DS Tmax values that can be offered over time for various split options and DS spectrum sizes. (SG size = 64, IPTV video channels)

Observe that the above figures do not propose a specific set of US or DS Tmax values that the MSOs should offer. An MSO with envisioned support for a particular Tmax value (depending on their specific demand and competition) can use the curves to estimate the lifespan of their network given that envisioned Tmax value.

Besides indicating the Tmax capabilities resulting from US split changes, node splits, DS spectrum size changes, and video delivery strategy changes, the above family of curves can also be used concurrently with other tools like the selective subscriber migration tool. For example, looking at Figure 4 which assumes digital video and SG size of 250 subscribers, it can be noticed that high-split with 1.2 GHz DS network cannot support DS Tmax values of 5 Gbps. However, if selective subscriber migration is used, where the few subscribers requiring 5 Gbps service are moved to another platform (like 10G PON over FTTH), then the rest of the subscribers on the network will need less demanding DS Tmax values like, say 1 Gbps, and therefore the same network with no modification will be able to support a DS service rate of 1 Gbps to 256 subscribers until 2026.

Finally, it should be noted that Figure 4 - Figure 9 include Tmax curves that can be supported using the 1x2.5Gbps GPON platform. For the digital video scenarios, it is assumed that video is transported via a

separate wavelength and therefore all of the GPON capacity is available for data. On the other hand, the IPTV scenarios assume that IP video traffic is sent over the PON technology along with data and therefore both video and data share the total available capacity, which leads to reduced Tmax values. It can be seen from Figure 4 - Figure 9 that HFC networks can compete very well against GPON using the available migration tools. This observation can lead us to argue that it is wise to limit the competition scope of HFC networks to only compete against EPON/GPON and use selective subscriber migration for customers who require service rates that are comparable to those offered by 10G PON, when the percentage of subscriber requiring those rates is small. This can be a smart move that yields a cost-effective migration strategy, where HFC networks can live for a long life while meeting the demand and addressing non 10G PON competition.

Finally, it should be noted that all of the curves in Figure 4 - Figure 9 will stretch to the right if the CAGR percentages drop.

6. Example Network Migration Strategy

As mentioned earlier, there are many tools available to help the MSOs with their HFC network migration exercise. Some of the near-term tools are:

- Enable more DOCSIS 3.1 OFDMA for the US
- Enable more DOCSIS 3.1 OFDM for the DS
- Node split/segmentation
- Enabling Switched Digital Video
- Increasing Video Compression
- Video BW reclamation by moving to IPTV
- Increasing the US split
- Increasing the DS spectral range
- Selective subscriber migration

An optimal network migration strategy is obtained by applying a comprehensive decision process that considers all the available tools concurrently. The process is repeated every time a decision is to be made. In particular, the optimal decision changes as time moves on. This paper listed various tools that can be used concurrently along with a proposed decision tree and family of time-aware curves to help make the right decision when a migration step is needed.

Based on the above, an example migration strategy is given below:

1. Reframe the goal as ‘constrained’ network migration **{Selective Subscriber Migration}**
 - a. The result of constraining the process is avoiding the costs associated with unnecessary ‘network-wide’ upgrades when only a small percentage of the subscribers demand a service requiring such an upgrade. This is accomplished using the Selective Subscriber Migration tool, where a constraint is put on the maximum data rate that should to be supported using the HFC network. Any peak rate exceeding this maximum is to be offered using another platform like PON over FTTH. An example of a chosen max rate to be supported on the HFC network can be 2.5 Gbps. Rates beyond 2.5 Gbps are not carried by the HFC network. That is, HFC is to compete with EPON/GPON but not 10G PON, when the percentage of subscribers requiring rates in excess of 2.5 Gbps is very small. The result is moving the small percentage of subscribers with demand for very high peak rates to a different platform, which in turn relieves the pressure from the existing network and therefore elongating the life of the network by allowing it to serve

- the remaining ‘normal’ subscribers without massive upgrades. Note that when the MSO decides to offer a particular service that requires an upgrade to the majority of their subscribers, then a network-wide upgrade will be appropriate and justifiable.
- b. The next steps in this migration strategy example assume network-wide upgrades, where all super subscribers have already been moved to a different platform.
2. Continue reducing the SG size **{Node splits}**
 - a. Node splits and node segmentation will help in reducing the SG size, which will reduce the overall busy hour BW requirements and therefore elongate the life of existing networks.
 3. Move to 204 MHz US with 1.2 GHz DS **{High-split/1.2 GHz DS network architecture}**
 - a. With DOCSIS 3.1, 204 MHz US enables the offering of 1+ Gbps US peak rates. 1.2 GHz DS adds additional spectrum to accommodate the increased DS traffic demand.
 - b. Assuming no node split (i.e., SG size of 250 subscribers), high-split with 1.2 GHz DS can offer symmetrical 1 Gbps until at least 2026 (DS limited). With two nodes splits (i.e., SG of 64 subscribers), the same architecture can offer 1 Gbps symmetrical service until 2031/2032.
 4. Continue Reclaiming video BW **{Move video to IPTV}**
 - a. Moving video to IPTV can further elongate the life of a high-split 1.2 GHz DS HFC network (with SG size of 64 subs) by two years. That is, it can offer symmetrical 1 Gbps service until 2033/2034.
 5. Further Increase the US & DS throughputs **{Move to either ESD or Dynamic Soft-FDD}**
 - a. DOCSIS 4.0 ESD specifications enables the US to go up to 684 MHz and the DS to go up to 1.8 GHz. This will yield US and DS peak rates of 5 Gbps and 10 Gbps, respectively. ESD is a fixed-split FDD operation that is easy to manage but it requires the replacement of taps and amplifiers. Amplifiers may be challenged with high gain requirements to accommodate the increased attenuation at higher frequencies. As mentioned earlier, limited TCP values may not a major issue for short and medium HFC plants.
 - b. Dynamic Soft-FDD is the FDX flavor for N+x networks. It enables the US to go up to 684 MHz and the DS to continue to be at 1.2 GHz, where the US & DS time-share the 108-684 MHz spectrum. This will yield US and DS peak rates of 5 Gbps and 10 Gbps, respectively. While Dynamic Soft-FDD does not require tap replacements, its operation may be a little more challenging than ESD. Also, due to the potential complexity of the amplifiers, it may only work for small cascades although it has the advantage of using the same FDX CPE silicon.
 - c. While ESD & Dynamic Soft-FDD are equivalent from capacities point view, the decision to go one direction versus the other has major implications to the network architecture and its operation.
 6. **Active Taps**
 - a. Taking the fiber deeper in the network beyond N+0 can be done in stages. In particular, the MSO can choose to maximize the use of existing hardlines by using active taps, which are small active amplification devices, to enable signal delivery at high frequencies over hardlines.
 7. **FTTT**
 - a. Whether active taps were used or not, before going FTTH, it is beneficial to visit this step to maximize the life of HFC networks by using the existing drop cables that can go beyond 25 GHz. In fact, with FTTT architecture, the HFC network becomes a point-to-point network with only the drop cable between subscriber and the network. This can enable FDX, ESD, or combination of the two.

8. FTTH

- a. The desired end goal of HFC networks is to extend the fiber to the home.

7. Conclusions

There are many tools available in the network migration toolbox which will enable the HFC networks to support their customers well into the 2030 decade. These tools include selective subscriber migration, node splits, moving video to IPTV, upgrading the US split, increasing the DS spectrum, dynamic Soft-FDD, active taps, FTTT, and FTTH.

After utilizing the selective subscriber migration concept to constrain the network migration process, the article proposed a decision tree combined with a time-aware decision process to yield an optimal network migration strategy. It was found that the combination of light/medium touch options like moving the US split 204 MHz with 1.2 GHz DS & nodes splits/segmentations can extend the life of the HFC network to 2030 if the desired DS Tmax is 2 Gbps or less. Deploying IPTV will further extend the life of HFC networks. Moreover, this article also showed that Dynamic-Soft-FDD & ESD can both increase the life-span of HFC networks. In the ESD case, going with US splits higher than 396 MHz when the DS is limited to 1.8 GHz is not optimal.

Abbreviations

bps	bit per second
BW	bandwidth
COVID-19	Corona Virus Disease 2019
DAA	distributed access architecture
DOCSIS	data over cable service interface specifications
DS	downstream
DSG	digital set-top gateway
EC	echo cancellation
EPON	Ethernet PON
ESD	extended spectrum DOCSIS
FD	fiber deep
FDD	frequency division duplex
FDX	full duplex DOCSIS
FTTH	fiber to the home
FTTT	fiber to the tap
Gbps	gigabit per second
GW	gateway
HD	High Definition
HFC	hybrid fiber coax
Hz	hertz
IG	interference group
IP	internet protocol
IPTV	Internet Protocol Television
kHz	kilo hertz
Mbps	mega bit per second
MHz	mega hertz
MPEG	Moving Picture Expert Group

MSO	multiple service operator
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
OOB	out of band
PON	Passive Optical Network
QAM	quadrature amplitude modulation
QoE	quality of experience
RF	radio frequency
RPHY	remote physical layer or remote PHY
SC-QAM	single-carrier QAM
SG	service group
SLA	service level agreement
STB	set-top box
TCP	total composite power
TG	transmission group
US	upstream
GPON	Gigabit PON
UHD	Ultra High Definition
UHS	Ultra-High Split
VOD	Video on Demand
XGS-PON	10G symmetrical PON defined by ITU

Bibliography & References

[ANGA-Cable-Migration-2019]	<i>“Navigating the HFC Network Migration Maze into the 2020 Decade”</i> , by A. Al-Banna, ANGA COM 2019.
[NCTA-COVID-19-Tracker-2020]	https://www.ncta.com/COVIDdashboard
[SCTE-Tec-Soft-FDD-2019]	<i>“Operational Considerations & Configurations for FDX & Soft-FDX: A Network Migration Guide To Converge The Cable Industry”</i> , by A. Al-Banna et. al., SCTE Cable-Tec Expo 2019.
[Spring-Forum-ESD-2016]	<i>“Using DOCSIS to Meet the Larger BW Demand of the 2020 Decade and Beyond”</i> , by T. Cloonan et. al., Spring Technical Forum, 2016.
[Spring-Forum-Migration-2016]	<i>“Network Migration Demystified in the DOCSIS® 3.1 Era and Beyond”</i> , by A. Al-Banna et. al., Spring Technical Forum, 2016.

DOCSIS® PKI: A Proposal for a Next-Generation Quantum-Resistant Infrastructure

Using Composite Crypto for Post-Quantum PKI Transitioning

A Technical Paper prepared for SCTE•ISBE by

Massimiliano Pala
Principal Security Architect
CableLabs
858 Coal Creek Cir., Louisville, CO 80027
+1 (303) 661-3334
m.pala@cablelabs.com

Table of Contents

1. Introduction.....	3
2. Trust Infrastructures and the Quantum Threat.....	3
2.1. PKI Building Blocks	6
2.2. Modern Cryptography and the Quantum Threat	7
2.2.1. Public-Key Cryptography	7
2.2.2. Hashing Algorithms	7
2.2.3. Encryption Algorithms	7
2.3. Algorithm Agility to the Rescue	8
3. A Composite Crypto-Based Solution.....	8
3.1. A New Paradigm: CompositeKeys and CompositeSignatures	9
3.1.1. Composed Public Keys and X.509 Certificates	10
3.1.2. Composite Signatures and X.509 Certificates	11
3.1.3. Generating Composite Signatures.....	12
3.1.4. Verifying Composite Signatures and Time-Dependent Validation Policies Deployment.....	12
3.1.5. Use of Composite Crypto for Backward Compatibility	13
4. A Complete Solution: From Requests to Revocation.....	13
4.1. Requesting Composite Certificates.....	13
4.2. Use of Composite Signatures in CRLs	14
4.3. Use of Composite Signatures in OCSP Requests and Responses.....	14
5. Composite Cryptography and Hardware Integration.....	14
6. Deploying a Backward-Compatible, Quantum-Safe DOCSIS PKI.....	15
6.1. Deploying Post-Quantum Solutions for RSA-Only Capable Devices	18
6.1.1. Fielded Devices and Authentications.....	18
7. Conclusion.....	20
Abbreviations	21

List of Figures

Title	Page Number
Figure 1 - Example of Composite Crypto Usage in X.509 Certificates	9
Figure 2 - The DOCSIS "New PKI" Hierarchy.....	16
Figure 3 - Validation Scenario Showing Multiple Entities with Different Capabilities.....	17

1. Introduction

The broadband industry has been relying on public-key cryptography (PKC) to provide secure and strong authentication across its networks and devices. In particular, the DOCSIS standard [Doc31, Doc40] uses X.509 [Itu509] certificates to verify that a device is a legitimate entity that is authorized to join the network—for example, a cable modem or a Remote PHY (R-PHY) node [Rphy1]. The choice of using digital certificates and public-key infrastructures (PKIs) to protect DOCSIS identities has resulted in a scalable and easy-to-deploy key management system for the entire industry.

Although the DOCSIS PKI has been a success story over the past 20 years (it is one of the largest PKIs ever deployed worldwide), things are changing rapidly on both the security side and the broadband industry side.

On the security side, new advancements in traditional and non-traditional computing are threatening our ability to use traditional public-key and key-exchange (KEX) algorithms. On the network infrastructure side, new zero-trust architectures are being designed that require software and hardware entities to securely authenticate to each other (and encrypt traffic) in a distributed environment.

This paper describes our proposal for a backward-compatible quantum-resistant trust infrastructure (or PKI) for the broadband industry. Specifically, our work focuses on the practical aspects of deploying a quantum-resistant trust infrastructure by leveraging our idea—namely, the composite cryptography mechanism [Com20].

The paper is organized as follows: Section 2 provides a description of the quantum threat for the various parts of a PKI; Section 3 describes the composite crypto solution and its two building blocks (i.e., `CompositeKey` and `CompositeSignature`); Section 4 describes how to practically deploy composite crypto in PKIs; Section 5 provides considerations surrounding the use of secure elements and hardware security modules (HSMs); and Section 6 describes a deployment proposal for securing the DOCSIS PKI. Finally, Section 7 provides our conclusions and envisioned future work.

2. Trust Infrastructures and the Quantum Threat

Deploying real security is difficult, and security engineers rely on basic cryptographic primitives to make sure protocols operate as intended and data is accessed only by authorized entities. Deploying real security is even more difficult when uncertainty around the efficacy of your basic tools is at risk. Unfortunately, we are living in a period of great cryptographic “uncertainty,” where the advancements in both traditional and quantum computing pose serious threats that may impact the possibility to provide secure access and data privacy not only for the broadband industry, but across the Internet.

Recently, advancements in quantum computing suggest that the possibility to break PKC might be closer than initially thought. Specifically, the work of Craig Gidney and Martin Ekerå [GE19] improves the efficiency of quantum computers to perform code-breaking calculations, thus reducing the required resources by orders of magnitude. For example, in 2015, researchers estimated that a billion quantum bits or qubits (aka q-bits) would be required to factor 2048-bit keys due to the need to use noise-reduction codes that require significant extra qubits themselves. With the recent advancement from Gidney and Ekerå, the number of required qubits is reduced to only 20 million.

Another example of the fast-paced rhythm of innovation occurring in the quantum space is the fact that researchers have already moved away from studying qubits to building quantum logic gates. Many

experts theorize that at the current speed of innovation, we will have a quantum computer within 20 years—a very short period of time in the cryptography space.

The takeaway message is alarming for all of us: Governments, military and security organizations, banks, medical facilities and everybody else will have no options to secure their data against powerful quantum computers.

The situation is even more alarming when considering that there are no complete or general solutions today that allow for securing all aspects of a PKI—not just signatures or certificates. In particular, regardless of the specific standard or technology used in a PKI (e.g., X.509, PGP), no deployments currently cover the possibility of using multiple algorithms in a combined fashion to further secure the infrastructure in case one or more algorithms are deemed compromised. Our approach addresses this limitation by leveraging a recursive construct for both public-key signatures and keys that can be applied to every aspect of the PKI lifecycle management.

On the quantum-safe cryptography standardization front, things are moving forward as NIST recently announced the beginning of Round 3 in its selection of a quantum-resistant cryptography standard that began few years ago [Nist20].

Because the new standard will specify one or more quantum-resistant algorithms for digital signatures, public-key encryption and key generation, the selection of finalists comprises various classes of algorithms and mathematical properties. In particular, of the original 69 proposals submitted for Round 1, only 26 made it to Round 2. For Round 3, only 8 candidates were selected. Four of the candidate cryptosystems provide public-key encryption, while the remaining four are digital signatures schemes.

The selected key encapsulation mechanisms (KEMs) that provide public-key encryption are:

- **Classic McEliece.** This KEM is the result of a merger between the Classic McEliece and NTS-KEM. This work is a code-based KEM based on the original cryptosystem from 1978. The cryptosystem has never gained much interest in the cryptographic community; however, its resistance to Shor’s algorithm makes it a good candidate for post-quantum standardization. Classic McEliece, which uses the binary Goppa code, comes with very large public keys but the smallest ciphertext of any other solution. Although this performance profile might not be the best fit for the general protocols we use over the Internet today, its stable specifications combined with a long history of cryptanalysis make it an appealing choice for some applications.
- **CRYSTALS-Kyber.** This algorithm is based on the presumed hardness of the module learning with errors (MLWE) problem. The scheme has excellent all-around performance for most applications. It also enables relatively straightforward adjustment of the performance-versus-security tradeoff by varying module rank and noise parameters. NIST regards this scheme as one of the most promising KEMs for standardization.
- **NTRU.** This is a structured lattice-based KEM that has an established adoption history because variations of this KEM have already been standardized by IEEE [Ntru09] and ANSI [Ntru10] organizations. Although NTRU has a small performance gap in comparison with Kyber and SABER, its longer history was an important factor in NIST’s decision because of less risk of unexpected intellectual property claims. An important characteristic of NTRU is the fact that it relies on a different problem than MLWE; thus, it provides diversity in the set of structured lattice-based KEMs for the finalists.

- **SABER.** This KEM is based on a variation of MLWE—namely, the Module Learning With Rounding (MLWR), in which rounding from one modulus to a smaller second one replaces the addition of small errors. In general, SABER provides good performance for general-purpose applications. One of the areas that NIST encourages more research on is side-channel analysis for the non-Number Theoretic Transform (non-NTT) style of multiplication that is unique to SABER. Together with NTRU and Kyber, SABER is one of the most promising KEMs selected for Round 3.

For digital signatures, the following are considered as finalists:

- **CRYSTALS-Dilithium.** This is a lattice-based signature scheme that relies on the MLWE hardness and the module short integer solution (MSIS). One of the advantages of CRYSTALS-Dilithium over its main competitor, Falcon, is the simpler implementation due to the use of the same modulus and ring for all parameter sets and samples. Overall, Dilithium has a good balance in terms of key and signature sizes and performs well for key generation, signing and verification, making it one of the strongest candidates for standardization as it performs well in real-world experiments.
- **Falcon.** The Falcon signature scheme is lattice-based and uses the “hash and sign” paradigm. Although this scheme is more complex to implement than Dilithium, it offers the smallest public key and signature sizes. From a key generation point of view, Falcon is slower than other candidates, but the overall strong performance makes it a good fit for existing Internet protocols and applications.
- **Rainbow.** This scheme is very different from the previous two. Rainbow is a multivariate signature scheme with an unbalanced oil and vinegar (UOV) construct. Rainbow provides fast signing and verification, along with very short signatures. The downside of this scheme is the extremely large public key sizes. Some issues with the security claims and the fact that NIST prefers algorithms with royalty free licensing put this algorithm at the bottom of the list.
- **GeMSS.** This is the second multivariate signature scheme that uses the “big field” paradigm. Although it offers the smallest signatures of any other schemes and provides a fast verification algorithm, the large size of public keys is the limiting factor for adoption, especially in low-end devices where signing can be very slow. Also, the large size of public keys makes it impractical when used with TLS or SSH without protocol changes. This algorithm is still in consideration, in case developments in Round 3 show the Rainbow scheme to not be suitable for standardization.

NIST also selected eight alternative algorithms whose evaluation will continue after the first selection for the standard. These eight alternative algorithms either might need more time to mature or are tailored to more specific applications. For example, the Sphinx+ scheme is considered very secure but also very conservative in its design. Because this translates into higher bandwidth and slower performance than the selected primary candidates, Sphinx+ is being considered only as a backup option for standardization. The review process for Round 4 will continue after Round 3 ends, and eventually some of these alternative algorithms could become part of the standard at a later date. The status of the NIST selection process is available in the NIST Internal Report 8309 [NISTIR].

2.1. PKI Building Blocks

To convey which parts of a trust infrastructure will be affected by the use of quantum computing, we provide a classification of algorithms or “building blocks” that are used to protect both authentication and user data and show how they might be impacted by the quantum threat.

Specifically, when it comes to the various types of building blocks, we shall differentiate between public-key algorithms for authentication (e.g., RSA, ECDSA), KEX algorithms (e.g., Diffie-Hellman, Elliptic-Curves Diffie-Hellman), hashing algorithms (e.g., SHA-256, SHA-3) and encryption algorithms (e.g., AES).

Public-Key Algorithms. Public-key algorithms are primarily used to authenticate (and sometimes encrypt) data. Algorithms like RSA or ECDSA are the most common when used in conjunction with X.509 digital certificates. The main difference between RSA and ECDSA, besides the underlying mathematical properties, is in the size of the cryptographic overhead (or bandwidth) and key-generation complexity. In particular, when compared through the performance lens, ECDSA has a clear advantage, especially when deployed in small or computationally limited devices. When compared through the security lens, though, further considerations are due. One interesting feature of RSA is its ability to use different key lengths without changing the algorithm. This allows, for example, RSA cryptosystems to increase the size of the public/private keypair to adjust to increased security risks due to advancements in computing and crypto research. ECDSA, instead, does not support this feature: Once the curve is selected (e.g., NIST’s Secp256r1 curve), the key size cannot be changed, and to increase the security of the system, new curves (e.g., NIST’s Secp521r1 curve) must be supported (even for validation only).

Key-Exchange Algorithms. This class of algorithms has been recently revamped because of the effort, in the TLS space, to deploy perfect forward secrecy (PFS). Specifically, the use of finite-field Diffie-Hellman (DH) and Elliptic-Curve Diffie-Hellman (ECDH) has been required by the latest TLS specifications [RFC 8446] to overcome the security limitations of the RSA-based KEX mechanism and decouple authentications from key exchanges. KEX algorithms are at the core of protocols like TLS to securely derive the encryption keys used after the negotiation phase.

Hashing Algorithms. This class of algorithms is at the core of the authentication process in modern PKIs. In fact, to authenticate any type of data (e.g., a certificate or simply a document), the data must be “summarized” to make the signing process efficient and more generic (i.e., providing the same operational approach when working with different algorithms). Because of this, breaking the properties of the hashing algorithms can be quite devastating for a public-key cryptosystem.

Encryption Algorithms. When it comes to securing data from unauthorized access, encryption algorithms provide the necessary building blocks via the use of securely exchanged (via the authentication process) encryption keys. The ability to successfully attack these algorithms can bypass any authentication process used during the transfer of encryption keys and therefore grant an attacker direct access to the data.

In the rest of this section, we provide a summary of the threats for each class of algorithms and how (and if) the quantum threat is significant in that space.

2.2. Modern Cryptography and the Quantum Threat

2.2.1. Public-Key Cryptography

When it comes to public-key algorithms such as RSA or ECDSA, the threat of quantum computers being able to “guess” (factor large numbers or solve the discrete logarithm problem) private keys is comparable for both cryptosystems. In fact, PKC uses mathematical algorithms to generate complex keys, thus making the code to reverse-engineer the private component from the public one statistically very hard. Different public-key systems can use different algorithms, as long as they are based on mathematical problems that are easy to put into place but difficult to reverse-engineer. For instance, any computer can multiply two extremely large prime numbers, but factoring the result is nearly impossible—at least, it would be for a classical machine. Although only few classes of algorithms (so far) have been identified to be more performant on quantum computers, factoring large numbers is one of them. Specifically, Shor’s quantum factorization algorithm could be easily used to break this type of cryptosystem.

The optimization that is possible on a quantum computer is due to the fact that it should be able to use the properties of quantum mechanics (e.g., entanglement) to probe for patterns within a huge number without having to examine every digit in that number. Because cracking both RSA and EC ciphers actually involves finding patterns in huge numbers, quantum computers can perform the inverse operation at practically the same speed as the forward one. For example, while on a conventional computer, finding a pattern for an EC cipher would take $2^{N/2}$ steps—where N is the number of bits in the key. On a quantum computer, the number of steps would be in the order of only $N/2$!

RSA, ECC and DH are examples of cryptosystems that will not be secure against an adversary with access to a quantum computer.

2.2.2. Hashing Algorithms

Hashing algorithms give us a chance to breathe a little easier. As explained in the previous section, some problems with a special structure (e.g., factorization for RSA, discrete log for EC) typically fall in the category of problems for which quantum computers can reduce the task to finding the period of some function, which is surprisingly not difficult to solve on a quantum computer.

However, for unstructured problems, quantum computers seem to provide “only” some non-trivial quadratic speedup, via the use of Grover’s algorithm [Gro96]. Simply put, this indicates that a hash function with 256 bits of security today would still have 128 bits of security in a post-quantum computing era.

2.2.3. Encryption Algorithms

Symmetric encryption, and more specifically AES-256, is believed to be quantum-resistant. Quantum computers are not expected to be able to reduce the attack time enough to be effective if the key sizes are large enough. Also in this case, the speedup from Grover’s algorithm allows quantum computers to perform exhaustive searches in the square root of the classical time, rendering classical attacks more expensive than generic ones in most cases. To achieve 128 bits of security in a post-quantum computing scenario (equivalent to AES-128 today), primitives providing 128 bits of security (e.g., AES) need to have a key length of at least 256 bits. In other words, the speedup in the quantum exhausting search provided by Grover’s algorithm is the main reason why the current recommendations for encryption with post-quantum computing security mandate for AES [Dr99] with a 256-bit key.

If you are using AES in your systems in 2020, you should favor AES-256 over AES-128. This is true for software environments, but it is especially important when chipsets and hardware deployment is involved; to make sure these devices can take you over the “quantum hump,” support for AES-256 is required [Aes16, Aes20].

Ultimately, as with the hashing algorithms case, modern symmetric encryption algorithms are not affected as badly by quantum computing as public-key ones as long as we deploy larger keys (e.g., AES-256).

2.3. Algorithm Agility to the Rescue

Algorithm agility is at the core of best practices when it comes to cryptosystems today. Algorithm agility refers to the possibility of substituting cryptographic algorithms (and associated data structures) as needed, without having to change protocol messages or main data structures.

Fortunately, the X.509 standard, used across the Internet and in the broadband industry, provides the possibility to extend the generic data structure to integrate new algorithm-dependent ones in order to identify public keys and signatures. Technically, this is achieved by coupling the crypto data structures with an algorithm identifier that can correctly validate and process the associated data structure(s) across different types of objects used in PKIs (e.g., X.509 certificates, Online Certificate Status Protocol [OCSP] responses, Certificate Revocation Lists [CRLs]).

For example, look at how public keys are encoded in X.509 certificates:

```
AlgorithmIdentifier ::= SEQUENCE {
    algorithm          OBJECT IDENTIFIER,
    parameters        ANY DEFINED BY algorithm OPTIONAL }

SubjectPublicKeyInfo ::= SEQUENCE {
    algorithm          AlgorithmIdentifier,
    subjectPublicKey   BIT STRING }
```

Here, `SubjectPublicKeyInfo` comprises an `AlgorithmIdentifier` that identifies the cryptographic algorithm and associated parameters, as well as a `subjectPublicKey`, which is a `BIT STRING` [RFC 5280]. The value of `subjectPublicKey` is the Distinguishing Encoding Rule (DER) encoding of the public-key structure as defined for the specific algorithm used.

For example, Section 2.3.1 of [RFC 3279] defines the contents of `SubjectPublicKeyInfo` and how to encode the `RSAPublicKey` structure whose DER representation is to be used for the value of `subjectPublicKey`.

Our solution leverages the algorithm agility feature and defines a new algorithm identifier that encodes, recursively, a set of one or more `subjectPublicKey` data structures to encode multiple keys with different algorithms, as discussed in the next section.

3. A Composite Crypto-Based Solution

The cable industry, as well as the larger Internet community, is faced with the very difficult task of addressing the quantum threat early in order to continue to securely authenticate network devices and users by switching to different algorithms when needed. On top of that, the broadband industry faces the additional challenge of handling this future transition while relying on fielded devices (e.g., CMs, R-PHY/R-MACPHY nodes) that might be expensive to replace or update.

Although the selection for standard post-quantum algorithms has not been finalized yet, we have been actively looking at how to deploy such algorithms for the broadband industry when they become available. In particular, we looked at how to protect the integrity of the most vulnerable parts of our DOCSIS trust infrastructure first (from a quantum-threat perspective)—that is, the root and intermediate CA. We focused on these assets first because they are valuable targets that would allow for an attacker to generate new valid entities for the ecosystem. Once the integrity of the upper levels of the hierarchy are secured, devices can be updated to support the deployed post-quantum algorithm(s) via Secure Software Download (SSD) or other mechanisms.

What we found in our research is that we could use the same algorithm agility feature that is built into modern PKIs to support the use of multiple keys for every aspect of the PKI lifecycle; from certificates to revocation lists, everything supports our new paradigm. In other words, our work enables the use of classic algorithms like RSA or ECDSA alongside new ones. This allows for a gradual transition to new algorithms without losing backward compatibility for devices that cannot be updated with the new algorithms. Figure 1 provides an intuitive representation of the composite crypto when used in X.509 certificates.

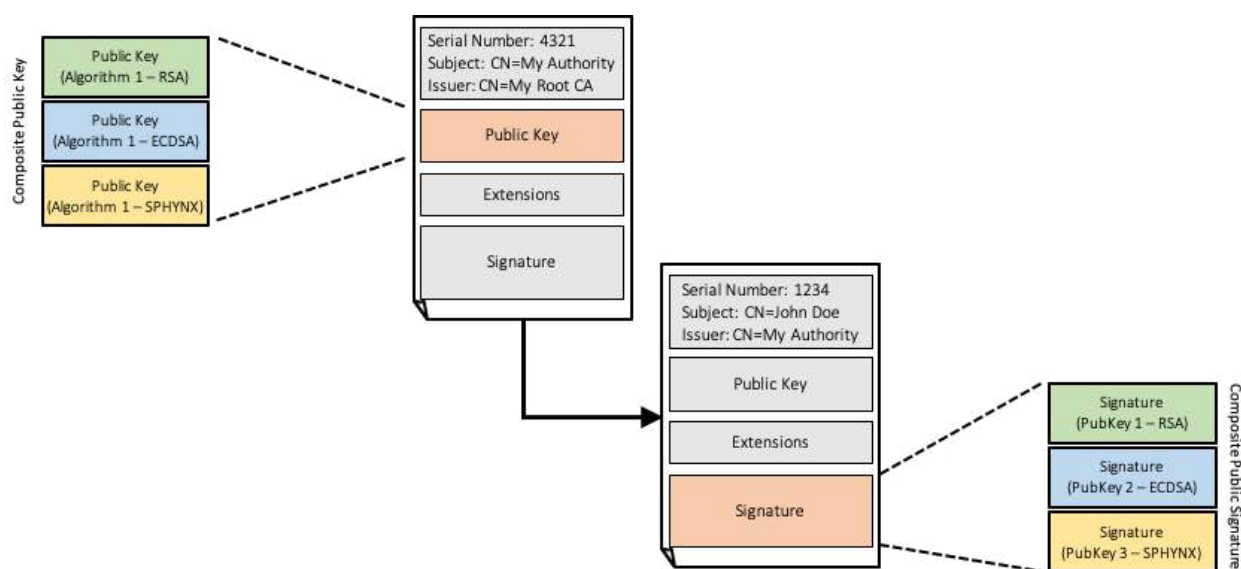


Figure 1 - Example of Composite Crypto Usage in X.509 Certificates

It is interesting to notice how our invention can be used at any time when there is either (a) the need to transition to a new protocol without establishing a completely separated infrastructure or (b) uncertainty related to the security of an algorithm over an extended period of time. For example, the solution described in this paper could be used to transition from less efficient cryptosystems (e.g., RSA) to more efficient ones (e.g., ECDSA) or to provide signatures with different hashing algorithms.

3.1. A New Paradigm: CompositeKeys and CompositeSignatures

As suggested earlier, to address our problem we relied on our initial considerations about algorithm agility to provide a simple and backward-compatible solution. In particular, we defined a new algorithm identifier and associated encoding that uses standard substructures to encapsulate multiple keys or multiple signatures in PKI data structures and authentication data.

The new type of public keys and signatures—namely, `CompositeKeys` and `CompositeSignatures`, respectively—provide the building blocks we were looking for: backward-compatible encoding for both signatures and public keys that allows for multiple keys and algorithms to be used to secure X.509 objects and produce generic signatures. By mixing classic and post-quantum algorithms, not only can all aspects of a PKI be protected today, but clients supporting at least one of the combined algorithms will be able to operate in the environment.

For example, to trust the authentication of data, relying parties might decide to verify one, some or all of the signatures depending on the ability of the relying party to support any of the algorithms used for keys and signatures.

3.1.1. *Composed Public Keys and X.509 Certificates*

The technical aspects of our work are simple. We first defined a new value for the `algorithm` field within the `AlgorithmIdentifier` used in the `SubjectPublicKeyInfo` of a `tbsCertificate` structure of a X.509 certificate:

```
compositeKeys OBJECT IDENTIFIER ::= {iso(1) identified-organization(3) dod(6)
                                     internet(1) private(4) enterprise(1) OpenCA(18227) 10 }
```

When this value is used for the algorithm identifier, it means that the value encoded in the associated public key field (e.g., the `subjectPublicKey` field) contains multiple public keys and associated parameters. The `parameters` field of the `AlgorithmIdentifier` itself shall be set to `NULL` in this case as there are no specific parameters associated with the composite key – the recursive nature of the data structure allows us to delegate the parameters to the individual keys definitions.

To encode the different keys, we use a sequence of `SubjectPublicKeyInfo` objects. Each of these objects encodes the specific algorithm identifier for the specific key with its parameters and value. The final sequence is then encoded as the DER representation of the sequence of keys.

We also define the `CompositePublicKeyInfo` as a `SEQUENCE OF SubjectPublicKeyInfo` where each `SubjectPublicKeyInfo` carries the information of one public key. The ASN.1 definition of the `CompositePublicKeyInfo` is as follows:

```
CompositeSubjectPublicKeyInfo ::= SEQUENCE (1..MAX) OF SubjectPublicKeyInfo
```

where the `SubjectPublicKeyInfo` within the `CompositeSubjectPublicKeyInfo` must not use `compositeKeys` as an algorithm identifier to prevent multiple levels of recursion.

For example, to add two separate public keys in an X.509 certificate via composite crypto, the encoding would be as follows:

```
aCompositeSubjectPublicKeyInfo = SEQUENCE { keyInfoOne, keyInfoTwo };
-- The main structure, a sequence of two subjectPublicKeyInfo

keyInfoOne.algorithm.algorithm = rsaEncryption;
keyInfoOne.algorithm.parameters = NULL;
keyInfoOne.subjectPublicKey = RSAPublicKey;
-- The keyInfoOne provides the definition for the first key (RSA)

keyInfoTwo.algorithm.algorithm = id-ecPublicKey;
keyInfoTwo.algorithm.parameters = EcpcParameters;
```

```

keyInfoTwo.subjectPublicKey      = ECPoint;
-- The keyInfoTwo provides the definition for the second key (ECDSA)

aCertificate.tbsCertificate.subjectPublicKeyInfo.algorithm.algorithm = compositeKey;
aCertificate.tbsCertificate.subjectPublicKeyInfo.algorithm.params = NULL;
aCertificate.tbsCertificate.subjectPublicKeyInfo.subjectPublicKey =
    DER(aCompositeSubjectPublicKeyInfo);

```

where `aCompositeSubjectPublicKeyInfo` is the sequence of two `subjectPublicKeyInfo` (i.e., `keyInfoOne` and `keyInfoTwo`). The DER representation of the `aCompositeSubjectPublicKeyInfo` is then stored in the `subjectPublicKey` field of the `subjectPublicKeyInfo` of the `tbsCertificate`.

3.1.2. Composite Signatures and X.509 Certificates

When it comes to signatures in X.509 certificates and their validation, we used a similar approach. We first defined a new algorithm identifier for `compositeSignatures` and then defined the specific data structures for the composite algorithm.

Specifically, when a `compositeSignatures` schema is used to encode multiple signatures at once, the value for the algorithm identifier associated with the signature is defined as follows:

```

compositeSignatures OBJECT IDENTIFIER ::= {iso(1) identified-organization(3)
    dod(6) internet(1) private(4) enterprise(1) OpenCA(18227) 11 }

```

When the `compositeSignatures` identifier is used, the corresponding value encoded in the `signatureValue` field contains multiple signatures and associated parameters encoded as the DER representation of a `CompositeSignatureValue` that is a SEQUENCE OF `SignatureInfo`. Each `SignatureInfo` carries the information about one of the signatures applied to the certificate. The definition of the `CompositeSignaturesValue` is as follows:

```

CompositeSignaturesValue ::= SEQUENCE (1..MAX) OF CompositeSignatureInfo

```

For example, to encode signatures made with two separate keys (one RSA key and one EC key), the encoding would be as follows:

```

aCompositeSignatureInfo      = { sigInfoOne, sigInfoTwo };
-- The main structure, a sequence of two SignatureInfo

sigInfoOne.algorithm.algorithm = rsaEncryption;
sigInfoOne.algorithm.parameters = NULL;
sigInfoOne.subjectPublicKey    = <RSA Signature Value>;
-- The sigInfoOne provides the definition for the first signature (RSA)

sigInfoTwo.algorithm.algorithm = id-ecPublicKey;
sigInfoTwo.algorithm.parameters = EcpkParameters;
sigInfoTwo.subjectPublicKey    = <ECDSA Signature Value>;
-- The sigInfoTwo provides the definition for the second signature (ECDSA)

aCertificate.signatureAlgorithm.algorithm.algorithm = compositeSignatures;
aCertificate.signatureAlgorithm.algorithm.params    = NULL;
aCertificate.signatureValue = DER(aCompositeSignatureInfo);
-- The final encoding of multiple signatures in a certificate

```

where the `aCompositeSignatureInfo` structure contains the sequence of the two `SignatureInfo` (i.e., `sigInfoOne` and `sigInfoTwo`). The DER representation of the `aCompositeSignatureInfo` is then used for the `signatureValue` field of the certificate structure.

3.1.3. Generating Composite Signatures

To generate composite signatures, the signer shall generate each signature independently by using each of the keys present in the signer's `CompositePublicKeyInfo` in the same order they appear. Specifically, the signer shall use the first key to generate the first signature, the second key to generate the second signature, and so on. The signer shall generate one signature for each key in the key set.

For example, if the `CompositeSubjectKeyInfo` has three public keys (K_1 , K_2 and K_3) of types RSA, EC and DSA, respectively, the signing party shall generate the first signature by using K_1 , the second signature by using K_2 , and the last signature by using K_3 .

3.1.4. Verifying Composite Signatures and Time-Dependent Validation Policies Deployment

To be able to verify composite signatures, a relying party shall verify each of the applied signatures independently. Also in this case, the relying party shall verify the signature by using the corresponding public key in the signer's certificate in order—that is, the order of the signatures within the `CompositeSignature` shall respect the order of the keys in the `CompositePublicKeyInfo` in the certificate.

For example, if the certificate has a `CompositeSubjectPublicKeyInfo` that contains three keys (K_1 , K_2 and K_3) of types RSA, EC and DSA, respectively, the relying party shall verify the first signature in the composite signature by using K_1 , the second signature by using K_2 , and the last signature by using K_3 .

One important aspect of our invention is that it can be combined with the possibility of applying validation policies that can be changed over time or remain static.

In a static configuration, for example, the relying party might set its policy not to evaluate the correctness of signatures if they do not support any of the used (or specific) algorithms, or otherwise refuse to trust the signed data entirely, even if it is not able to verify just one of the composite signature's elements.

In a time-dependent policy, relying parties could instead use the quantum threat risk level to set the threshold for the policy change. Imagine, for example, an infrastructure in which both RSA and ECDSA algorithms are used via `CompositeKeys` encoded in the root and intermediate CA certificates. Also imagine that the RSA algorithm is set to retire because deemed not secure in 10 years (e.g., the used key sizes for the infrastructure are not considered secure anymore). In addition, assume that ECDSA will still be considered secure for the application (e.g., larger keys can be deployed here because of better performances). A validation policy could allow relying parties to validate composite signatures by using only the RSA algorithm for the next 10 years. After that, the policy might mandate relying parties to validate signatures by using all algorithms to make sure the stronger one(s) is validated too.

In the post-quantum scenario, this translates to a very similar approach.

Specifically, although static policies might be more appropriate for those relying parties or devices whose crypto cannot be updated (see the next section), more dynamic ones could be deployed when the validation of new algorithms can be added to the device. In this case, up to a certain security risk level

(e.g., until practical deployment of quantum computers is achieved), relying parties and devices could still be allowed to use just the traditional algorithms for validation and enable the new ones when the risk level for the involved stakeholders goes over the acceptable threshold (and support for it is successfully deployed).

3.1.5. Use of Composite Crypto for Backward Compatibility

The same solution can be used when deploying a new infrastructure where participants in the ecosystems might not be able to update their security parameters. In this case, composite crypto structures can be used to deploy trust infrastructures where new and old algorithms coexist in the `CompositeKeys` and `CompositeSignatures` of certificates. The deployment of superseded algorithms along with new ones allows relying parties that cannot update their cryptographic suites (e.g., devices that are already in the field) to participate in the same infrastructure while still allowing other relying parties to use stronger validation algorithms.

In other words, composite crypto can be used to keep using old algorithms to accommodate for older, already-in-the-field devices that might have hardware constraints (e.g., they have a secure element that cannot be replaced) without compromising the overall security of the infrastructure. The stronger keys/algorithms will be used by more capable devices (e.g., P-521 or quantum-resistant algorithms).

4. A Complete Solution: From Requests to Revocation

Although other solutions have been tried to provide support for adding multiple keys or algorithms to certificates by adding new types of extensions, no other solution actually tackles all the aspects of the PKI lifecycle. Specifically, no other solution is available that addresses not only the authentication of certificates but also the authentication of certificate requests and revocation objects. In this section, we take a look at the applicability of our solution to these aspects that are central to the correct behavior of PKIs.

4.1. Requesting Composite Certificates

In PKIs, the [PKCS10] standard is commonly used when it comes to requesting certificates. Many standard (and non-standard) protocols use it as the core building block for their own certificate request messages. Examples of this can be found in Certificate Management over CMS (CMC) [RFC 5272, RFC 5273] and in the Automated Certificate Management Environment (ACME) [RFC 8555].

Fortunately, our work is compatible with the PKCS#10 format.

In particular, to authenticate PKCS#10 requests with composite crypto, the `signatureAlgorithm`'s algorithm identifier in the `CertificationRequest` structure can be set to carry the `compositeSignatures` value, and the `parameters` one can be set to `NULL`.

The `signature` field is the one that carries the DER representation of `CompositeSignatures` and contains all the signatures generated with the `compositeKeys` associated with the identity that is requesting a certificate. The signatures are calculated, as usual, over the DER representation of the `certificationRequestInfo` field of the `CertificationRequest`.

4.2. Use of Composite Signatures in CRLs

CRLs are the oldest form of revocation for X.509 certificates [RFC 5280, RFC 5759, RFC 6818]. Their structure was inspired by credit-card number blacklists and are used to convey the list of serial numbers of certificates that have been revoked (together with an optional reason code). Because this list is often signed by the issuing CA (or a designated signer), we need to make sure that this list is securely authenticated, even in a post-quantum threat scenario.

Our approach seamlessly works with CRLs too.

As with the case of X.509 certificates, the `signatureAlgorithm` field in the `CertificateList` structure can be set to carry the `compositeSignatures` value, and the `parameters` field can be set to `NULL`. The `signature` field of the `CertificateList` can be set to carry the DER representation of the `CompositeSignaturesValue`.

Also, in this case, there is no change in how the signatures are generated because the individual signatures are calculated over the DER representation of the `tbsCertList`, as usual.

4.3. Use of Composite Signatures in OCSP Requests and Responses

OCSP requests and responses have signature fields that can be leveraged with composite signatures to address the quantum threat without requiring any protocol changes.

To authenticate OCSP requests, the `signatureAlgorithm` algorithm identifier in the `Signature` structure of the `OCSPRequest` can be set to `compositeSignatures`, and the `parameters` field can be set to `NULL`. The corresponding `signature` field of the `Signature` structure can then hold the DER representation of the `CompositeSignature` value itself. The signatures are calculated, as usual, over the DER representation of the `tbsRequest` in the `OCSPRequest` structure.

For OCSP responses, the `BasicOCSPResponse` structure provides, together with the `tbsResponseData` and the `signatureAlgorithm` ones, the needed fields to host composite signatures. Specifically, the `signatureAlgorithm` algorithm identifier in the `BasicOCSPResponse` structure can be set to `compositeSignatures`, and the `parameters` field can be set to `NULL`. The corresponding `signature` field can be set to hold DER representation of the `CompositeSignaturesValue`. Also in this case, the individual signatures can be calculated and encoded, as usual, over the DER representation of the `tbsResponseData` field of the `BasicOCSPResponse`.

5. Composite Cryptography and Hardware Integration

Modern cryptography relies on two main principles: making algorithms public and moving all security properties to the secrecy of keys. That is why one of the pillars of modern cryptography is keeping your secrets ... secret! This is true not only for shared secrets but also for private keys.

To help with the security of keys, best practices commonly require the use of HSMs or secure elements in our devices to make sure that (a) private key computations safely happen on a dedicated processor and (b) their value cannot be extracted by a remote party.

From this point of view, the quantum threat not only poses a risk from a security perspective, but it also requires hardware updates to provide the same security properties that we have enjoyed for decades.

Although quantum-resistant algorithms must be run in software until support for them is provided via secure hardware implementations, the composite cryptography solution itself is fully compatible with existing hardware and security modules. This is because the processing of the individual keys and signatures still relies on the same primitives; therefore, no changes are required for composite crypto integration, as long as the algorithm is supported by the crypto accelerator. Similarly, current standard interfaces to crypto hardware, like PKCS#11 [PKCS11] and supporting libraries, do not require any changes for the same reasons. This is extremely important for preserving backward compatibility with deployed HSMs, which are usually large and very expensive pieces of equipment used mostly to secure operating CA keys.

This means that all investments that Certificate Service Providers (CSPs) made in purchasing and maintaining their HSMs are not impacted, as no changes are needed to leverage composite crypto. For example, existing root and intermediate CAs can add new keys to their own certificate and have their new request signed by using composite crypto today. This allows current infrastructures to add new algorithms and still be able to leverage the security (and certifications) of today's crypto hardware (e.g., FIPS 140-2).

Summarizing, composite crypto can be used today with existing certified hardware components, thus allowing the transition from, for example, RSA to ECDSA without the need for new hardware or certifications. Support for new algorithms is still required for deploying quantum-safe crypto.

6. Deploying a Backward-Compatible, Quantum-Safe DOCSIS PKI

So far, the DOCSIS ecosystem has deployed two different infrastructures throughout its lifetime. The first one, today referred to as “legacy PKI,” was deployed 20 years ago and provides its services for DOCSIS 1.1–3.0 devices. Because this infrastructure is set to expire soon, it was reasonable to focus our efforts on the newer infrastructure.

The second deployed infrastructure, or “new PKI,” was introduced to update the security parameters for the whole broadband industry and provides its services to secure not only DOCSIS 3.1–4.0 devices but also other entities associated with the existing and upcoming distributed architectures (e.g., R-PHY or CCap Core). Because this infrastructure is not going to sunset anytime soon, our work has been focused on defining the deployment strategy for securing this second “new” infrastructure across the “quantum-threat hump.”

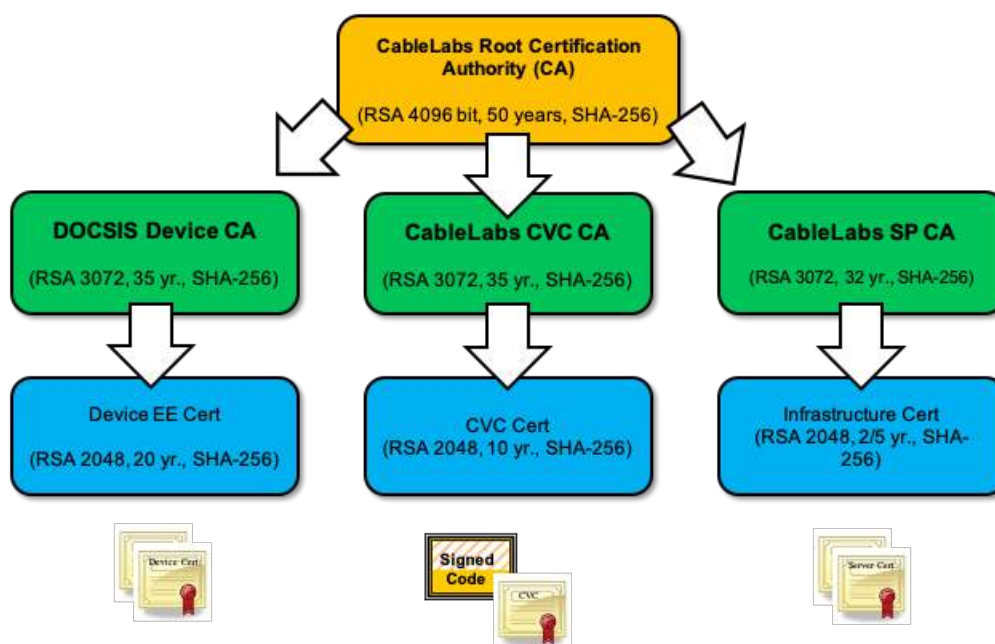


Figure 2 - The DOCSIS "New PKI" Hierarchy

Figure 2 depicts the “new PKI” structure that comprises a three-tier hierarchy with an offline root CA that issues a second level of intermediate CAs. These CAs issue only end-entity certificates and have assigned operational scopes (e.g., device certificates vs. code-signing certificates) that limit their liability in case of compromise. All participating entities in the infrastructure use the RSA algorithm for their keys.

Currently, our work is focused on setting up a test infrastructure that uses composite keys and signatures to secure the core part of the infrastructure (i.e., the root CA and the intermediate CAs) against the quantum threat and, at the same time, provide the possibility to leverage algorithms other than RSA for increased efficiency.

The envisioned test infrastructure mimics the current “new PKI” hierarchy and uses, for the core of the hierarchy (i.e., root and intermediate CAs), three separate public keys: the current RSA key, a new ECDSA key and a new Post-Quantum-Algorithm (PQA) one. Device or end-entity certificates are issued with a single algorithm that can be either RSA, ECDSA or the PQA, depending on the entity or device capabilities. Network or server-side identities are issued with composite crypto keys comprising all the deployed algorithms to support all classes of devices (i.e., classic and post-quantum).

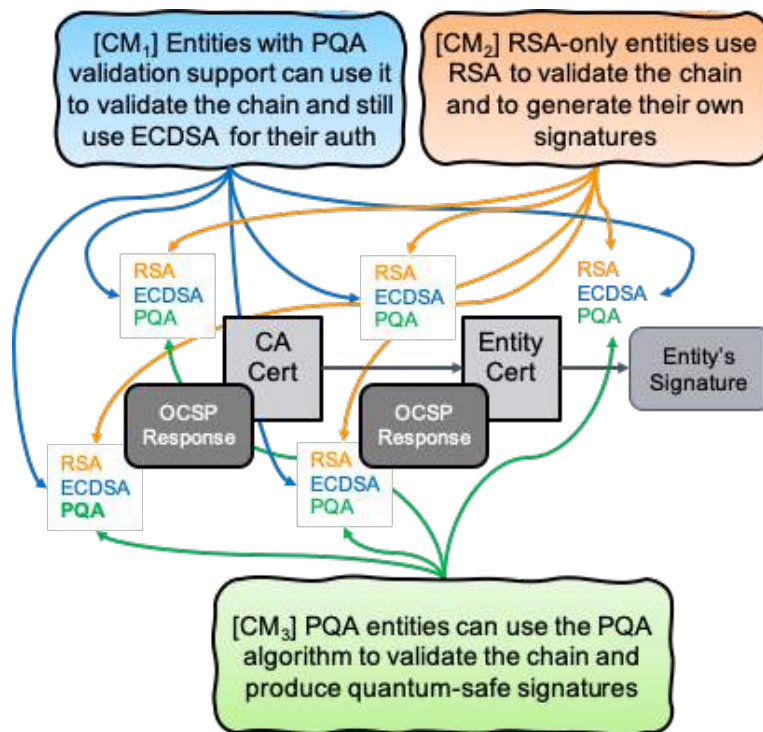


Figure 3 - Validation Scenario Showing Multiple Entities with Different Capabilities

Figure 3 provides a validation scenario in which three devices with different capabilities are authenticating themselves with credentials from the quantum-safe PKI. In this scenario, CM₁ is capable of only working with ECDSA P-256 keys but can validate PQA signatures and public keys, CM₂ is capable of only working with RSA keys, and CM₃ fully supports PQA (not only validation but also signing and key management) and can also validate RSA and ECDSA.

On the client side, CM₁ can benefit from using either classic or quantum-safe crypto and still authenticate itself by using classic crypto (ECDSA). CM₂, instead, is a constrained device and can only use RSA; without any further update, this class of devices cannot be securely authenticated under the quantum-threat model (see the next section for proposed mitigations to address this use case). CM₃ is a newer device that fully supports the selected PQA and therefore uses that algorithm to both validate the network credentials and generate its own authentication traces. In case CM₃ also supports classic cryptography, RSA or ECDSA validation algorithms can still be used before quantum-based attacks become practical.

On the network side, the CMTS must support all the algorithms supported by the entire population of deployed devices—both for validation and authentication. To be able to authenticate CM₁, the CMTS must support ECDSA. To be able to authenticate CM₂, instead, the CMTS has to support RSA. Ultimately, authenticating CM₃ requires the CMTS to support the PQA. On the authentication side, to allow devices that might support only one of the deployed algorithms to be able validate the network credentials (e.g., in Baseline Privacy Plus Interface [BPI+] V2), the CMTS must support all of the used algorithms in its composite key. This allows the use of RSA, ECDSA and PQA on the various classes of devices without the need of separate identities or infrastructures.

In the rest of this section, we provide an overview of the proposed approaches for tackling the quantum threat when deploying full PQA-based solutions is not an option—for example, because of limitations in fielded devices.

6.1. Deploying Post-Quantum Solutions for RSA-Only Capable Devices

One of the most challenging issues when it comes to cryptography is to include devices with different capabilities, and some of these devices might not be upgradeable. This can be due to software limitations (e.g., firmware or applications cannot be securely updated) or hardware limitations (e.g., crypto accelerators or secure elements).

Although composite cryptography cannot be used to secure non-quantum-safe authentications against the quantum threat by itself, in case entities and devices do not support any post-quantum algorithm, we identified a solution that can be used to extend the lifetime of deployed devices for the broadband industry.

6.1.1. Fielded Devices and Authentications

To consider securing fielded devices that cannot be updated to support PQAs, we looked at their current capabilities. We researched which classes of quantum-safe algorithms are available and how can we leverage them, given today's hardware constraints, to provide quantum-safe authentications.

What we found is that especially for constrained devices, the only option at our disposal is the use of pre-shared keys (PSKs) to allow for post-quantum safe authentications for the various identified use cases. We looked at the limitations and how to tackle them when planning for the transitioning. In this scenario, the post-quantum PSK (PQP) is used to generate quantum-safe signatures and, in some cases, to also provide a “second factor” of authentication for the certificate chain when no PQA support is available.

To generate quantum-safe signatures, devices can start using the PSK with their device private keys (e.g., RSA keys) to produce quantum-safe authentication data; the classic cryptography provides the identity information, along with the proof of possession of the classic private key, while the PQP provides the security of the message via a symmetric signature. Combining the PSK with the authentication process can be done in different ways. A hash-based key derivation function (HKDF) [RFC 5869] can be used with the PQP to derive a message-specific key that is then used with an HMAC function to authenticate the messages.

In BPI+ Version 2¹, because of the use of the Cryptographic Message Syntax (CMS) [RFC 5652], combining PSKs with DOCSIS authentication could also be used to provide key encapsulation capabilities for delivering authorization keys via a quantum-safe mechanism [RFC 8696]. For previous versions of DOCSIS, or where BPI+ V2 is not supported because direct RSA encryption of the authorization key is used, additional changes to the protocol messages might also be required.

Let's now take a look at two different scenarios and how to address their limitations. The first use case assumes that private keys, certificates and crypto capabilities cannot be updated—in this case, we rely on traditional crypto to perform the needed setup operations securely before the quantum threat is real. To remove the requirement for time-safe deployment, a second proposal is introduced that looks at devices whose private keys cannot be updated, but their support for composite crypto and quantum-safe KEX algorithm can (e.g., via SSD).

¹ The new version of BPI+ was recently introduced in the DOCSIS 4.0 specifications to introduce several improvements to the authorization process, such as mutual authentication and perfect forward secrecy.

6.1.1.1. *Immutable Devices and Quantum-Safe Traditional Authentications*

In this scenario, we look at entities and devices whose support for new algorithms cannot be updated—not even for validation-only operations (i.e., no support for private keys or signing). Because of these restrictions, our proposal is to leverage traditional cryptography to distribute per-device PSKs that can be leveraged for post-quantum authentications.

Specifically, our proposal is to enhance the DOCSIS protocol to introduce the possibility to securely transfer (or derive) a common PSK between the operator's network and the device being authenticated (i.e., the cable modem or the R-PHY node). Once the PQP is securely delivered to the device, this secret can stay dormant until needed for generating quantum-safe signatures. This PSK can be deployed as part of the initial registration of devices to the network, or it could be initiated at any time as long as the PQP is transferred securely.

One very important aspect of the solution is to make sure not only that the session parameters are properly authenticated via the quantum-safe signature, but that the certificate chain is protected against modifications by including it into the original signature. Assuming the PSK is secure, the relying party (e.g., the CCap Core or the CMTS) can trust both the signature and identity of the device because of the security of the PQP.

The big limitation here is related to the security of the PSK. Because the PSK has to be transferred or derived by using traditional cryptography, an attacker could potentially pre-record the device's traffic and then—when access to a quantum computer is obtained—get access to the PQP by breaking the classic KEX algorithm. The attacker would then be able to impersonate any device, even when using the PQP when generating signatures. Although a more secure solution is provided in the next section, operators can make things difficult for a malicious attacker who is pre-recording DOCSIS traffic to analyze and decrypt it at a later time.

Indeed, operators can deploy keys at random intervals or use procedures for combining new and old values (and/or replace them) at random times. An attacker would require knowledge of the whole history of the device connectivity to be able to attack its PQP.

6.1.1.2. *Partially Upgradeable Devices and Quantum-Safe Traditional Authentications*

When entities and devices can be updated to support new algorithms, but their private keys cannot (e.g., they are tied to secure elements that cannot be updated), more secure options can be adopted for transferring the PSK. We think that this upgrade path might be the most common for the broadband industry given that the possibility to provide secure software updates is built into the DOCSIS protocol since its inception.

In this scenario, we assume that devices have been updated to support composite crypto and a quantum-safe KEX algorithm, but they cannot update their own private keys. We also assume that the root and the intermediate CAs have been deployed with a PQA algorithm alongside traditional ones.

To protect the PQP against a possible all-powerful adversary that can break the traditional cryptography, we share the PQP by using a quantum-safe KEX algorithm. The use of a quantum-safe KEX algorithm guarantees that an adversary would not be able to have access to the PQP, even when pre-recording encrypted sessions. The use of traditional cryptography still provides the needed secure identity validation to make sure the PQP is shared across the right entities. Also, in this case, when generating authentication

data, we need to protect the certificate chain because the link between the intermediate CAs and the end-entity certificate is not yet protected via a PQA.

When the entity's certificate can be updated and signed by using a PQA (i.e., the CA signs a new certificate for the entity that includes the original "traditional" key of the entity only, and it is signed using all the keys in the CA's composite certificate), the need to sign the certificate chain with the PSK can be relaxed. In fact, because the links from the root to the end entity is already secured by the use of a PQA, no additional use of the PSK to protect the device's or CAs' identities is needed.

7. Conclusion

In this paper, we have examined the quantum threat, the current status of post-quantum cryptography and the associated standardization efforts. We also describe how the quantum threat will affect the various classes of algorithms we use today within the broadband industry.

The core of this paper provided a description of our novel approach based on composite cryptography, and how it can be used to address the quantum threat. Specifically, we provided the technical description of the composite cryptography building blocks (`CompositeKeys` and `CompositeSignatures`) and showed how to integrate them in every aspect of modern trust infrastructures and associated services (e.g., certificates, CRLs, OCSP). We then explored our proposal for a quantum-resistant and backward-compatible DOCSIS PKI and the use of PSKs to secure entities and devices that will not have access to quantum-resistant cryptography.

Our future efforts will be aimed at working on open source tools and test environments. The quantum-safe test services deployment will be paramount for security experts and researchers to experiment with combining different quantum-resistant algorithms and study their interactions with our protocols. The selection of the protocols and their parameters will pave the road to quantum-resistant DOCSIS implementations and deployment.

Ultimately, the takeaway message from our work is that the quantum threat is closer than many people think, and we need to be already preparing our infrastructures and protocols for the upcoming revolution. Not only quantum computers will become more capable of handling more complex problems; the advancements in quantum-based algorithms put everybody's security and privacy at risk.

Although the deadline for planning to address this new class of threats is fast approaching, our work shows how the broadband industry can already start to address them today and lead the transition to quantum-safe cryptography to be able to continue to protect users' privacy and securely deliver top-quality services.

Abbreviations

AES	Advanced Encryption Standard
BPI+	Baseline Privacy Plus Interface
CA	certification authority
CCAP	Converged Cable Access Platform
CRL	certificate revocation list
CSP	certificate service provider
DER	Distinguished Encoding Rules
DOCSIS	Data Over Cable Service Interface Specifications
EC	Elliptic-Curves
ECC	Elliptic-Curves Cryptography
ECDH	Elliptic-Curves Diffie-Hellman
ECDSA	Elliptic-Curves Digital Signing Algorithm
EE	end entity
DH	Diffie-Hellman
HSM	hardware security module
IETF	Internet Engineering Task Force Standards Organization
ISBE	International Society of Broadband Experts
KEM	key encapsulation mechanism
KEX	key exchange (algorithm)
MLWE	module learning with errors
MLWR	module learning with rounding
MSIS	module short integer solutions
NIST	National Institute of Standards and Technologies
NTT	Number Theoretic Transform
PKC	public-key cryptography
PKCS#10	Public Key Cryptography Standard 10 (certificate request)
PKCS#11	Public Key Cryptography Standard 11 (hardware interface)
PKI	public-key infrastructure
OCSP	Online Certificate Status Protocol
PFS	perfect forward secrecy
PQA	post-quantum algorithm
QC	quantum computing
R-PHY	Remote RF Layer (PHY)
R-MACPHY	Remote Media Access Control and RF Layer (PHY)
RSA	Rivest-Shamir-Adleman (cryptosystem)
SHA-256	Secure Hash Algorithm 2 (256-bit)
SHA-3	Secure Hash Algorithm 3
SCTE	Society of Cable Telecommunications Engineers
SE	secure element
SSD	secure software download
TLS	Transport Layer Security
UOV	unbalanced olive-vinegar (construct)

Bibliography & References

Doc40: *Data-Over-Cable Service Interface Specifications, DOCSIS 4.0, Security Specifications*. CableLabs Publication, 2019. Available as CM-SP-SECv4.0-IO1-190815.

Doc31: *Data-Over-Cable Service Interface Specifications, DOCSIS 3.1, Security Specifications*. CableLabs Publication, 2020. Available as CM-SP-SECv3.1-IO9-200407.

Ntru9: *Institute of Electrical and Electronics Engineers (2009) IEEE Standard 1363.1-2008 – Specification for Public Key Cryptographic Techniques Based on Hard Problems over Lattices* (IEEE, Piscataway, New Jersey, United States). Available at <https://doi.org/10.1109/IEEESTD.2009.4800404>

Ntru10: *American National Standards Institute (2010) ANSI X9.98-2010 – Lattice-Based Polynomial Public Key Establishment Algorithm for the Financial Services Industry* (ANSI, New York City, United States). Available at <https://webstore.ansi.org/standards/ascx9/ansix9982010r2017>

Itu509: ITU-T Recommendation X.509 (2005) | ISO/IEC 9594-8:2005, *Information Technology - Open Systems Interconnection – The Directory: Public-key and attribute certificate frameworks*.

Rphy18: *Data-Over-Cable Service Interface Specifications, DCA – MHA v2. Remote PHY Specification*. Available as CM-SP-R-PHY-I10-180509.

Com20: M. Pala. *Composite Public Keys and Signatures*, IETF I-D 03. July 2018.

GE19: C. Gidney and Martin Ekerå, *How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits*, Available at <https://arxiv.org/abs/1905.09749>

Nist20: National Institute of Standards and Technology, *Post-Quantum Cryptography – Round 3 Submissions*. Available at <https://csrc.nist.gov/Projects/post-quantum-cryptography>

NISTIR: Gorjan Alagic et Al., *Status Report on the Second Round of the NIST Post-Quantum Cryptography Standardization Process*, NIST, July 2020. Available At <https://csrc.nist.gov/publications/detail/nistir/8309/final>

Gro96: Lov K. Grover. *A Fast Quantum Mechanical Algorithm for Database Search*. In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing*, Philadelphia, Pennsylvania, USA, May 22–24, 1996, pages 212–219. ACM, 1996.

Dr99: Joan Daemen and Vincent Rijmen. *AES proposal: Rijndael*. 1999.

Aes16: Markus Grassl, Brandon Langenberg, Martin Roetteler, and Rainer Steinwandt. *Applying Grover’s Algorithm to AES: Quantum Resource Estimates*. In *PQCrypto*, Volume 9606 of *Lecture Notes in Computer Science*, pages 29–43. Springer, 2016.

Aes20: Xavier Bonnetain, María Naya-Plasencia and André Schrottenloher. *Quantum Security Analysis of AES*. *IACR Transactions on Symmetric Cryptology* Vol. 0, No. 0, pp.1–3, 2020.

RFC 8696: IETF RFC 8696, R. Housley, *Using Pre-Shared Key (PSK) in the Cryptographic Message Syntax (CMS)*, December 2019.

RFC 8555: IETF RFC 8555, R. Barnes, et al., *Automatic Certificate Management Environment (ACME)*, March 2019.

RFC 8446: IETF RFC 8446, E. Rescorla, et al., *The Transport Layer Security (TLS) Protocol, Version 1.3*, August 2018.

RFC 6818: IETF RFC 6818, P. Yee, *Updates to the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*, January 2013.

RFC 6402: IETF RFC 6402, J. Schaad, *Certificate Management over CMS (CMC) Updates*, November 2011.

RFC 5869: IETF RFC 5869, H. Krawczyk and P. Eronen, *HMAC-based Extract-and-Expand Key Derivation Function (HKDF)*, May 2010.

RFC 5758: IETF RFC 5758, Q. Dang, et al., *Internet X.509 Public Key Infrastructure: Additional Algorithms and Identifiers for DSA and ECDSA*, January 2010.

RFC 5652: IETF RFC 5652, R. Housley, *Cryptographic Message Syntax (CMS)*, September 2009.

RFC 5280: IETF RFC 5280, W. Polk, et al., *Cryptographic Message Syntax (CMS)*, May 2008.

RFC 5273: IETF RFC 5273, J. Schaad, et al., *Certificate Management over CMS (CMC): Transport Protocols*, June 2008.

RFC 5272: IETF RFC 5272, J. Schaad, et al., *Certificate Management over CMS (CMC)*, June 2008.

RFC 3279: IETF RFC 3279, W. Polk, et al., *Algorithms and Identifiers for the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*, April 2002.

RFC 2986: IETF 2986, M. Nystrom, et al., *PKCS #10: Certification Request Syntax Specification*, Version 1.7, November 2000.

PKCS11: OASIS Standard, S. Gleeson and C. Zimman, *PKCS #11 Cryptographic Token Interface Base Specification*, Version 2.40, April 2015.

Live Single-Stream HDR Production for HDR and SDR Cable Distribution

What Happens When “Post Production” Is Real-Time

prepared for SCTE•ISBE by

Chris Seeger

Director, Advanced Content Production
NBC Universal, LLC
900 Sylvan Avenue
Englewood Cliffs, NJ 07632
chris_seeger@nbcuni.com

Table of Contents

Title	Page Number
1. INTRODUCTION	3
2. CONTENT CONVERSION AND PERCEPTUAL MEASUREMENT	3
2.1. <i>How the Human Visual System Sees Light and Color</i>	5
2.2. <i>SINGLE-STREAM HDR AND SDR PRODUCTION</i>	9
2.3. <i>Hybrid Log-Gamma</i>	11
2.4. <i>HLG from camera to display</i>	12
2.5. <i>Color Conversion and objective color accuracy measurements</i>	12
3. SINGLE-STREAM HDR-SDR PRODUCTION	15
4. HDR/SDR SINGLE-STREAM DISTRIBUTION	18
5. CONCLUSION	20
ABBREVIATIONS	21
BIBLIOGRAPHY & REFERENCES	21

List of Figures

Title	Page Number
Figure 1 - How Do We See Light and Color? (© Universal Pictures from PacRim)	4
Figure 2 - The Human Visual System and Color	6
Figure 3 - Static versus Dynamic Adjustments to Light Levels	6
Figure 4 - The Human Visual System is Relative	7
Figure 5 - Color Appearance and HDR Production	8
Figure 6 - Nothing in This Image is Actually Moving	9
Figure 7 - Single-Stream HDR and SDR Production (ITU-R Report BT.2408)	10
Figure 8 - Objective Color Accuracy Measurement	13
Figure 9 - LUT Conversion (devices & software): Sarnoff(SRI) "Yellow Brick Road" Pattern	14
Figure 10 - LUT Conversion: Delta-E ITP (SDR-to-HDR) Using Sarnoff(SRI) Color Checker 2014	15
Figure 11 -: Examples of Shading: No Highlights	17
Figure 12 - Examples of Shading: Highlights	18
Figure 13 - This conversion produces a PQ image which is visually identical to HLG	19
Figure 14 - A simplified view of the NBCU/Comcast production through distribution in PQ	19

1. Introduction

Developing video content for the extended dynamic range and color depth in high dynamic range (HDR) and wide color gamut (WCG) now common in consumer displays isn't necessarily new: the first functional 4K/HDR televisions came out in 2016, and popular over-the-top (OTT) video platforms have routinely produced content in HDR since then. However; it is far, far more difficult to shoot and deliver HDR content for live linear workflows than it is for most file-based post produced approaches. That's because in live HDR production environments, handling graphical placements, color correction, and luminance normalization across what is usually a mix of SDR (standard dynamic range, which is essentially contemporary HDTV) and HDR cameras must be done on-the-fly and in real-time.

Also, differences in system colorimetry exist within the primary HDR systems, specifically Hybrid Log-Gamma (HLG) and Perceptual Quantizers (PQ), and took us some time to understand. A functional workflow for live delivery requires multiple levels of conversion for various signals throughout the infrastructure - like graphics, video playback, SDR and native HDR cameras. This three-part paper will explore 1) conversion and perceptual measurement for HDR and SDR; 2) how HDR production workflows are orchestrated for live content, and 3) how those live linear workflows are distributed – a topic likely of most interest to the cable television industry.

In addition, the paper will illuminate the research and collaborative processes necessary to build a functional, single-stream HDR-SDR workflow, at each stage of the pipeline, without compromising the artistic intent or quality of the distribution paths for either the HDR or the primary (and revenue-generating) SDR. In-depth research of color conversions will be discussed and explored, such as the specialized techniques described in the International Telecommunications Union – Radiocommunications Sector's (ITU-R) Recommendation BT.2124. Production layouts will be explored and described, so that readers gain a deeper understanding of decisions made in the baseband (serial digital interface(SDI) and IP – Internet protocol), file-based conversions/transcodes, and orchestration layers.

2. Content Conversion and Perceptual Measurement

Content conversion and perceptual measurement matters because of the vast differences in a scene you see with your own eyes, and that same scene captured by a camera and displayed on a screen. In real life, we see these bright and colorful scenes – but when we capture and move the image through to legacy TV's, some of the highlights and light levels get muted. That's because a traditional TV can was designed to display 100 nits (a “nit” is a measure of brightness, roughly equivalent to the light of one candle per square meter). Figure 1 shows a simulated scene represented in both HDR (lower image) and SDR.




Figure 1 - How Do We See Light and Color? (© Universal Pictures from PacRim)

The interplay between cameras, display screens and brightness dates back to the early days of analog cathode ray tube (CRT), which could achieve only a relatively marginal level of brightness, measured at approximately 100 nits. Compare this against today's HDR TVs (and some smartphones) – which can be as bright as 4,000 nits; and SDR TVs, which is to say “traditional HDTVs,” that can scale the peak white to a maximum of around 400 nits (note that a recommendation that is in development would make it around 200nits so it matches reference/graphics white in HDR). As time went on, and digital techniques replaced analog, cameras advanced to the point of being able to record a much wider dynamic range – meaning more detail to describe a larger range of luminance, from dark to bright, and a far richer span of colors that is closer to what the human visual system is capable of. HDR(High Dynamic Range) gives you substantially better highlights (the twinkly stuff), and WCG(Wide Color Gamut) gives you redder reds, bluer blues, blacker blacks.

HDR got a better foothold, commercially, within the medical imaging community, which required a method to see in a highly detailed and reliable way. As it turned out, the HDR images were so vivid, and so much closer to the workings of the human visual system (which we'll get to next), that it began to be developed for consumers. In today's consumer electronics marketplace, it's getting harder to find a television that doesn't include HDR, than one that does.

Dolby Laboratories (Dolby), in particular, focused on a system of quantization that could standardize the way HDR content is displayed, known in the industry as “PQ,” for perceptual quantization. In essence, PQ adds more bits to areas where the human visual system can see more detail and fewer bits where we see less detail. The thought is, if we can't see the detail, why store more detail. Additionally, PQ decided to use an absolute mapping which associates a specific code value with a luminance level.

If history is our guide, a proposed video standard isn't really a technology story unless a competing standard is vying for consideration. Which, in the case of HDR video, is HLG, co-developed by the British Broadcasting Corporation (BBC) and Nippon Hoso Kyokai (NHK – also known as the Japan

Broadcasting Corporation). Unlike PQ, HLG deliberately has some aspects of backwards compatibility to SDR(SDR, which is to say “HDTV” that uses traditional gamma mapped signals) when used with wide color gamut or ITU-R Recommendation BT.2020. Backward compatibility does not exist with older legacy displays that use the narrower color space in HDTV (ITU-R Recommendation BT.709). It works by defining a nonlinear transfer function, known as an electro-optical transfer function (EOTF), in which the lower half of the signal values use a gamma curve, and the upper half use a logarithmic curve.

SDR and HLG are relative systems that scale the entire image from a lower luminance level to a higher luminance value (the dynamic range doesn’t change only the luminance level). SDR was designed as a 100 nit system and we scale it to 200-400 nits for consumer delivery.

HLG was originally designed as a 1000 nit system and was adapted with a newer algorithm (ITU-R Recommendation BT.2100) in order to scale to other luminance levels without hue shifts. Scaling shifts shadows, midtones and highlights which can be an important consideration when distributing a signal to consumer displays that are getting brighter and brighter. Scaling does not increase the dynamic range (it’s like turning the volume up in audio).

The differences between the absolute (PQ) and relative (HLG) HDR systems are important to remember when we get to the distribution section of this document.

Our work producing and distributing HDR content began after a test run with the 2015 Independence Day fireworks in New York city, and then in earnest with the opening ceremonies for the 2016 Olympic Games, in Rio de Janeiro, Brazil. Since then, Comcast/NBCUniversal produced and distributed HDR video for the 2018 PyeongChang Winter Olympics; the 2018 and 2019 Notre Dame football season; the 2018 men’s Federation Internationale de Football (FIFA) games; and, earlier this year, Chicago Blackhawks hockey games. We learn something new every time we venture into an HDR/WCG event, and in particular these learnings tend to focus on ways in which to optimize content conversions, and avoid unnecessary duplications. The aim is a live production that “produces once, outputs twice:” Once in native HDR; once in converted SDR. This is no small feat. The rest of this paper will explain why.

2.1. How the Human Visual System Sees Light and Color

If the ultimate goal of video production is to faithfully reproduce what the eye sees, then a rudimentary understanding of the human visual system (HVS) is useful. The basics: the human eye takes in light, which is processed primarily by cone cells, which are sensitive to certain wavelengths – long, medium and short. Our rod cells are sensitive to dark scenes, and thus aren’t as involved in the processes of, say, watching television. Rod cells are more numerous than cone cells, but are less important, perceptually, during normal lighting conditions.

Figure 2 illustrates how the HVS processes information. Incoming light is captured by the three photo receptors (cones) in the eye, which have peak sensitivities in the Long, Medium and Short wavelengths, designated “L,” “M” and “S.” Linear light is converted into a non-linear signal response, to mimic the adaptive cone response of the HVS. Notice the different ratio of L/M to S wavelengths.




Figure 2 - The Human Visual System and Color

Not only that, but the HVS is adaptive, so, the responses can change, and those changes must be taken into account, especially when considering a much wider dynamic range.

Figure 3 illustrates static versus dynamic adjustments to light levels. Adapting to darker environments – like attempting to see objects in the night sky, after being in a lit room – can take as long as 30 minutes; in brighter environments, it can take five to 10 minutes for the eyes to adapt. In general, the HVS can see about 12 simultaneous stops; “stops” are a scaled measure of dynamic range. But! The HVS can operate within a 24-stop range, by doing a number of things, like scanning up and down in different brightness ranges, then adapting(perceptual) while adjusting the human pupil (physical), which when dilated, allows us to see darker detail, but loses brighter detail. Pupil contraction allows us to see detail in brighter objects, but sacrifices darker detail.




Figure 3 - Static versus Dynamic Adjustments to Light Levels

Because of adaptive factors (some of which aren't mentioned), the HVS is relative. It plays tricks on you!

Allow me to prove it to you:

In Figure 4, the stripe in the middle of the image appears brighter on the left side than the right – but it's not. As you scan from left to right on Figure 4, the images will appear to change – but they are not. It's merely your brain, adapting. At any moment in time, the HVS is juggling a complex mixture of perception, adaptation, sensitivity, acuity, and “day vs. night” variables. As content producers, it's important to pay attention to these variables, so as to provide consistent content to our viewers.



Figure 4 - The Human Visual System is Relative

Then there's the matter of how the Human Visual System interprets color and contrast, and how many different ways the brain can see things.

For instance:

- The “Hunt Effect,” simulated in the upper portion of Figure 5, illustrates how color saturation appears to increase – but the only thing that actually increased was the luminance/brightness.
- The “Stevens Effect,” does something similar with contrast, in that it appears to increase with brighter light.
- The “Bartleson-Breneman Effect,” investigates how display image contrast increases with the luminance of surround lighting.
- The “Helmholtz-Kohlrausch,” shows that where there is a saturated and a neutral color of the same luminance, the saturated color may appear brighter. This is another example of how saturation plays a role in perceived brightness, which will change with the viewing environment.
- The Perkinje Effect (illustrated in the lower image in Figure 5) describes the tendency for the peak luminance sensitivity of the eye to shift toward the blue end of the color spectrum at low illumination levels as part of dark adaptation. In consequence, reds will appear darker relative to other colors as light levels decrease




Figure 5 - Color Appearance and HDR Production

Because of the vast increases in luminance that come with HDR technologies, a significant issue to resolve is overall viewing comfort. Simply put, with HDR, as objects of certain colors get larger, they can become uncomfortably bright to view. The ITU-R produced standards for all levels – but they don’t include object size, color and temporal elements. As a result, small, twinkly images can be really beautiful, but large, bright objects can quickly become offensively bright.

This requires broadcasters and video content producers to purposefully shade material by observing additional elements mentioned earlier (object size, color and duration on-screen). Eventually, my hope is that “color loudness” rules will be established, just as there are for audio. Although out of scope for this paper, we recently filed a preliminary patent with some concepts which we believe will help to develop a “comfort metric” for HDR measurement tools.

Figure 6 is a bit of “eye-candy” to expose other strange perceptual elements. In the image, nothing is moving, but the blue dots appear to be moving. It illustrates the impact of drastically different contrast levels, which essentially make our brains freak out.




Figure 6 - Nothing in This Image is Actually Moving

2.2. SINGLE-STREAM HDR AND SDR PRODUCTION

Figure 7 takes us directly into the deep waters for content producers. It depicts the inner workings of a video control facility, as well as a control room, based on ITU-R Report BT.2408. In essence, we mapped out all of the pieces of the live HDR puzzle, necessarily including cameras, displays, the video switcher, graphics, distribution outputs, and how a typical HDR production would be laid out. Few (if any) broadcasters have budgets to support a full complement of HDR cameras; generally speaking, the HDR cams are used for “important capture,” while the extant suite of SDR cameras are put to use in, say, helicopter shots or those less vital to the overall imagery. Over time, all the camera’s will be capable of HDR.

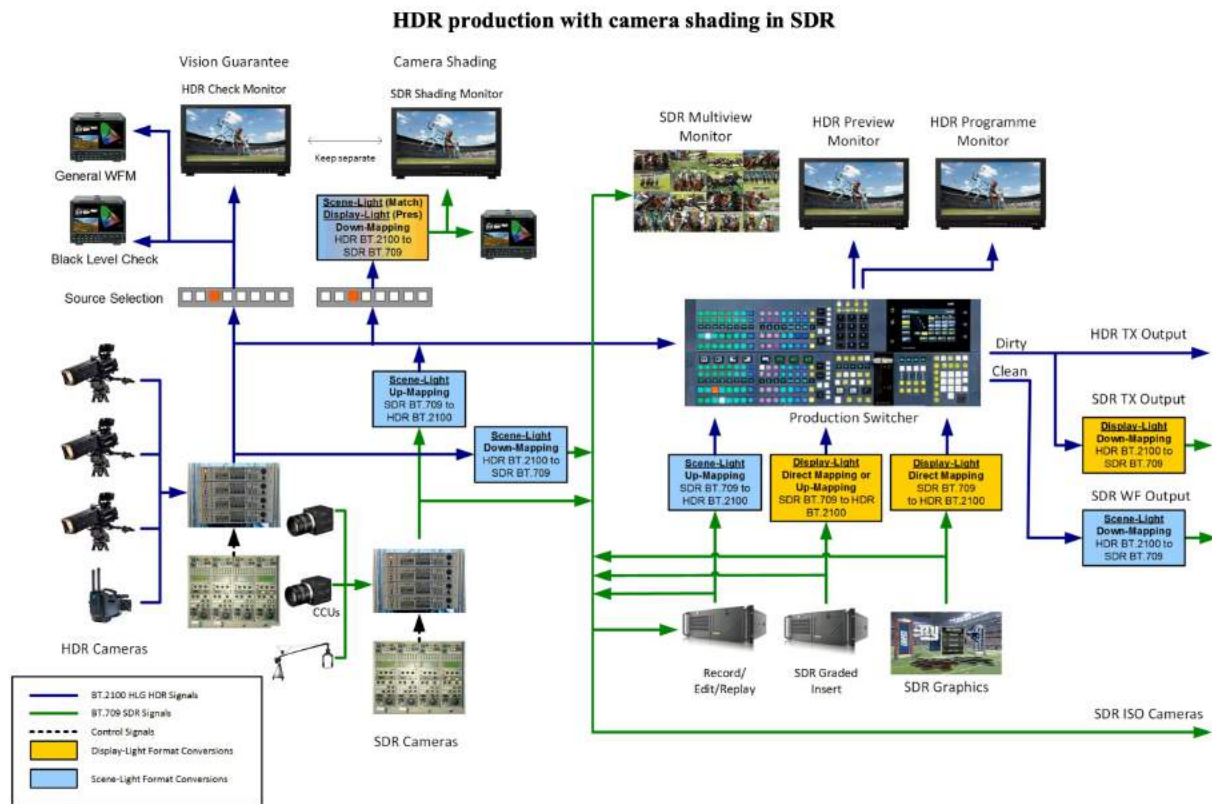


Figure 7 - Single-Stream HDR and SDR Production (ITU-R Report BT.2408)

2.2.1. Video Control

The video control area includes both HDR and SDR cameras, managed by a “shader” – essentially a “golden eye” for color and light – who matches the outputs of both camera types before they go on air. All sources need to be converted to the same HDR system before they enter the production, therefore, SDR sources are converted by “direct-mapping” or “up-mapping,” to HDR, using specific look-up tables (LUT). Specific LUTs are used and are chosen dependent on each production. If, for instance, the production uses S-log3 versus HLG there will be a different LUT set. Also, some productions prefer a different mapping of light levels requiring a different set of LUTs. The primary goal for the design of any LUT-set is to preserve artistic intent on the HDR and SDR sides of the output.

In HLG productions, two types of conversion must be applied. One is called “Scene Light Conversion,” which gets applied to older SDR camera’s in order to match them with native HLG’s “natural look”, and the other is “Display Light Conversion.” which get’s applied to all other inputs and outputs (HLG to SDR).

Scene Light Conversion applies the source video’s inverse-OETF (optical-to-electrical transfer function) to generate scene light prior to conversion, using the output OETF (to HLG, SDR or PQ) as the video

signal. In essence, it converts the SDR signal, with its color balance, and rebalances the colors in order to better match with a native HLG camera. HLG-BT.2100 cameras produce what has been labeled as the “Natural Look”. Some markets set their camera’s to output what has been labeled the “Traditional Look” which restores a similar color balance and saturation compared to SDR or PQ. This could complicate media exchange in the future.

SDR-to-HLG Display light conversion (DLC) applies a the SDR EOTF prior to the conversion to HLG, in order which converts the source to displayed light which properly preserves the artistic intent (color balance) of the source. SDR-to-HLG scene-light conversion will desaturate images while HLG-to-SDR scene-light conversion would oversaturate images because the variable luma-gamma in HLG will not have been applied.

Predictive LUTs are also an important video control room component, so that HDR to SDR conversion can be previewed, such that we know exactly what the SDR output feed to legacy distribution will look like, before it goes to air.

2.2.1. Control Room

In the control room, depicted in the right box of Figure 7, a multi-viewer screen is used by the director and technical director, to examine all camera outputs and select which shot to use. A LUT can be used here which will output the HDR feed into the SDR multi-viewer compilation. The control room is also where slow motion replays, SDR graphics, and HLG replays are composited in the video switcher, all while preserving the color balance of the source. Those elements are direct-mapped using Display Light Conversion (DLC). Native HDR playback and HLG replays do not require upconversion, and feed directly into the video switcher.

The switcher outputs directly to a native HDR distribution feed, which then goes to the transmission system; the signal will also be transmitted in SDR to our legacy broadcast network. Legacy SDR output from the HDR production requires a DLC (HDR-to-SDR), which reduces the dynamic range but maintains the color balance and therefore as much of the original artistic intent in the HDR production as possible. We have found that mapping HLG-75% to SDR-95% and then constructing a subjective knee which preserves additional highlights from HLG between SDR 95% and 109%. An HLG “pre-compression” knee applies compression natively in HLG from 90%-109% and then applies an additional knee which “post-compresses” HLG 75%-109% to SDR 95%-109%. This process creates substantially more compression in upper end of the HLG curve beyond SDR “legal range” and preserves additional highlights within SDR “legal range.”

2.3. Hybrid Log-Gamma

HLG works by capturing scene light, then sending it through an OETF. The result is a curve, representing the scene’s light level, that determines the quantization of the light levels into certain bits or code values. The scene light signal is what gets carried to the display, via SDI, IP or file for video distribution. Inside the receiving display, assuming it is HLG-equipped, several things happen: it applies an inverse-OETF, to reverse what was done in the camera, then an OETF is applied using gamma-to-luma only (not R,G,B - red, green, blue), so that it can scale the luminance of the images depending on the display’s peak

brightness capabilities. A 1,000 nit display will use one overall system gamma value, while a 2,000 nit version will use another, and that gamma value will scale only the light level, separate from the color. Perceptual effects like the Hunt Effect are important to remember when varying luminance and still require more study to understand what their impact might be. Variables like the perceptual effects are part of why we convert to PQ for distribution.

2.4. HLG from camera to display

HLG signals, from camera to display, apply a variable luma gamma in the display, which is determined by the display's peak brightness capability. By contrast, SDR and PQ apply gamma to R, G & B. This difference in gamma application for HLG produces a non-linear difference in chromaticity, relative to SDR or PQ.

The offshoot of this rather difficult process is that the signal looks desaturated compared to what we're used to in SDR. In some countries, that desaturated look is considered "less colorful," even though more accurate to the scene. The desaturated version is commonly known as a "natural look," whereas a "traditional look" re-saturates the images. As a direct result, HDR cameras carry two HLG settings – one for the natural look, the other for the traditional look. Deciding and/or converting between the two "looks" adds an additional level of complexity in broadcast workflows. A "Mild" setting is provided in some camera settings. It achieves a setting in between natural and traditional.

2.5. Color Conversion and objective color accuracy measurements

At this point, we know that we have to convert between SDR, HLG, or PQ. For the live-linear realtime conversions we commonly use 3D LUTs with a minimum size of 33 points and recommend a high quality interpolation method. 3D LUTS allow our file-based conversion workflows to match our baseband hardware conversions (live and post production conversions will match). Newer conversion methods are becoming available that allow for purely mathematical transforms. Purely mathematical transforms perform much less interpolation than LUTs and may become preferred over time.

In the early days of our HDR journey, we made subjective measurements for every LUT and every piece of new conversion hardware – we'd put the video in (test patterns and actual video), then examine the conversion vs the original side-by-side on a displays, and on a scope. It was a fairly inaccurate way of looking at very complex things like light and color. We found ourselves needing an objective measure that could make critical measurements. The ITU adopted metric created by Dolby (BT.2124) for objective color volume measurement. ITU-R BT.2124 uses a "color difference metric, called Delta-E ITP which measures the color volume difference (Delta) of one normalized sampled color volume (HLG, PQ and SDR) and another. ITP is a derivative of ICtCp which applies a small conversion to Ct ($Ct/2$). I and Cp remain the same.

ICtCp was designed to use key aspects of the HVS, in three steps, illustrated in Figure 8.




Figure 8 - Objective Color Accuracy Measurement

First, incoming light is captured by the three photo receptors (cones) in the eye, that have the previously discussed peak sensitivities in the Long, Medium and Short wavelengths. Next, the linear light is converted into a non-linear signal response, to mimic the adaptive cone response of the HVS. It is ICtCP's ability to algorithmically mimic the HVS that allows us to measure the perceptual appearance. It's not perfect, but it goes a long way to mimic how our eyes and brains translate color and light.

It's worth noting that this was important enough for us to build a plug-in to perform this measurement, using mathematical formula's contributed by Dolby and other recommendations from Philips. Normalizing all signals into ICtCp-PQ-BT.2020 (one large color volume container) enables us to make a single plot for comparisons of conversion accuracy between all current SDR-HDR-WCG systems.

First, after normalizing the signals, the T/P (color) components of ITP can be plotted on an X-Y axis to check for hue shifts throughout the LUTs conversions points.

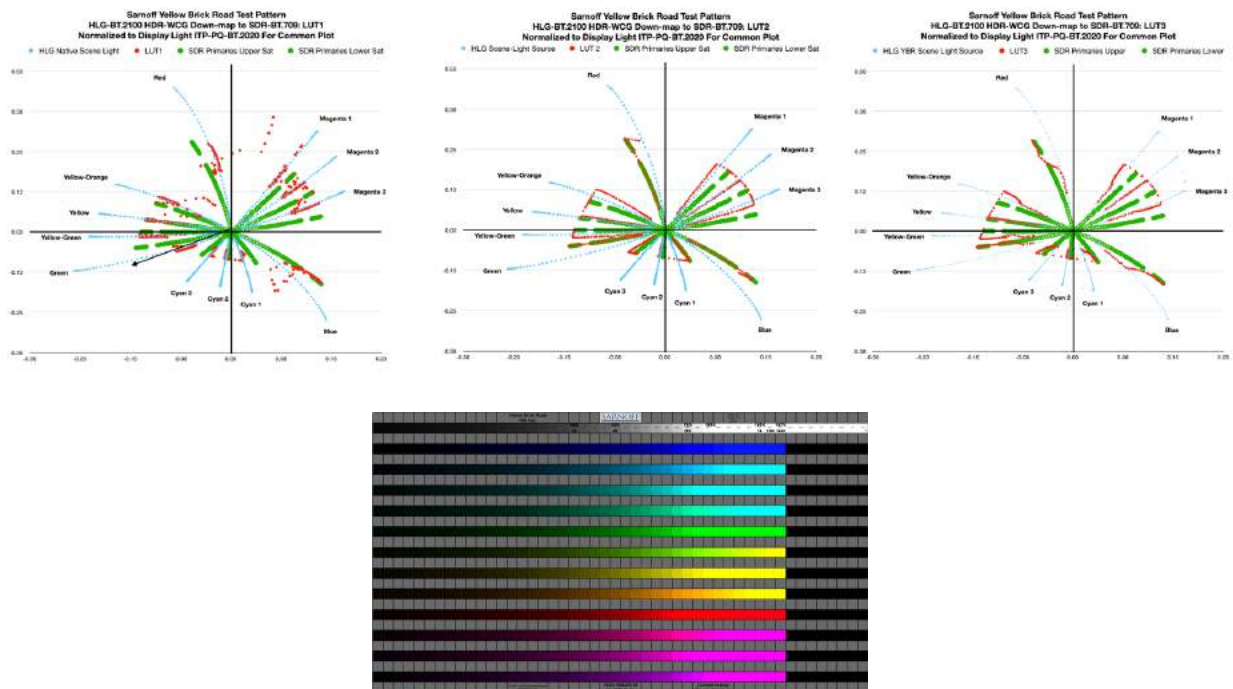


Figure 9 - LUT Conversion (devices & software): Sarnoff(SRI) "Yellow Brick Road" Pattern

Figure 9 compares three LUTs, with 576 sample points, to test up/down mapping and round-trips for color accuracy using the T/P components of ITP. All three samples are measuring the same pattern – a compilation of 576 color chips, to ascertain the ramp up in primary colors, and also the accuracy of compound colors (magenta, yellow, cyan, etc). In each sample, the blue dots represent the original(following BT.2020 primaries), the red dots represent the conversion, and. The green lines show the measurement of the SDR color primaries (BT.709) , because we want to know if the conversion is tracking the BT.2020 or BT.709 primaries so we can understand the conversion strategies.

LUT 1, on the far left, shows a conversion following the ITU-R Report BT.709 primaries for RGB for the majority of the mappings, but is extremely inaccurate with shadows and has less color detail. LUT 2, in the middle, shows excellent linear tracking of a conversion from ITU-R Recommendation BT.2020 to follow the BT.709 primaries. LUT 3, on the right, is a conversion that reproduces the BT.2020 colors within the BT.709 color triangle until the saturation/luma gets too high, then it slews toward the BT.709 primaries.

Of the three, the left-most conversion (LUT 1) is the worst of the bunch; the middle (LUT 2) is very good, and the far right is somewhat more subjective, if only because of the inevitable light and color tradeoffs that occur when converting from a larger to a smaller container (from HDR to SDR; BT.2020 to BT.709).

Figure 10 is an example of Delta-E ITP measurements for a conversion from SDR to HLG. Delta-E ITP provides a simple metric representing the difference of intensity and color between one image and another. An major benefit to the Delta-E ITP algorithm is that any unit measured above a “1” is considered a “Just Noticeable Difference,” or JND. This means that humans are able to discern a difference between two object’s light or color, if it measures above a 1. When calculated one against the other with the Delta-E ITP algorithm, it gives us a reliable indicator of noticeable differences, indicating accuracy issues with color and light conversions between HLG, PQ and SDR.

ITP Difference Original vs LUT Converters	
Colorchecker CSM Suite HLG 709 colors in 2020	Display Referred SDR->HLG
	$\Delta E\text{-ITP}$
dark skin	0.6107
light skin	0.4276
blue sky	0.5727
foliage	0.4997
blue flower	0.4987
bluish green	0.3500
orange	0.9263
purplish blue	0.4047
moderate red	0.8893
purple	0.3227
yellow green	0.5079
orange yellow	0.5393
blue	0.4771
green	0.5339
red	0.7252
yellow	0.3631
magenta	0.7047
cyan	0.3159
white	0.2364
Gray 2	0.2208
Gray 3	0.2589
Gray 4	0.2547
Gray 5	0.2193
Gray 6	0.1750

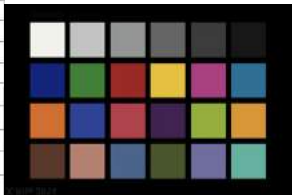


Figure 10 - LUT Conversion: Delta-E ITP (SDR-to-HDR) Using Sarnoff(SRI) Color Checker 2014

Since implementing the original color metrics workflow, we’ve also added the ability to plot our SDR, PQ, HLG sources normalized into CIE1976 u’v’ in order to examine absolute chromaticity based on recommendations from Philips in ITU.

3. Single-Stream HDR-SDR Production

The bulk of our experience in producing live content in HDR is, perhaps predictably, sports. As mentioned previously, our HDR/WCG adventures began in earnest with the 2016 Olympic Games, in Rio de Janeiro, Brazil. Subsequent HDR productions included the 2018 PyeongChang Winter Olympics; the 2018 and 2019 Notre Dame football season; the 2018 men’s FIFA games; and, earlier this year, Chicago Blackhawks hockey games. Producing live content in HDR, while preserving and/or converting live feeds for SDR broadcasts, involves a blend of baseband and file-based workflows and overall orchestration.

Specifically, three different tone maps are used to manage overall brightness, for scene light, HDR display light, and SDR display light. Live HDR cameras feed into the switcher are output in native HDR. Legacy SDR camera's are up-mapped using Scene-Light so that the SDR color balance matches HLG. Other legacy pre-produced content and graphics from SDR are direct-mapped using Display-Light on the fly to preserve artistic intent.

File movement is challenging, from an orchestration perspective. Files need to be moved efficiently, converted, and transcoded and then blended with the live, native feeds.

Consistently shading HDR content is another significant component of live productions, to set levels, color balance and obviate uncomfortable light levels. Through the ITU-R, broadcasters have established recommendations for a reference white / graphics white level, at 203 nits (75% HLG; 58% PQ in narrow range.) Reference white acts as an "HDR anchor point" for all other focal points like skin tones, grass, and other images, to prevent HDR brightness levels from going too high -- while reserving enough space for highlights up to 1,000 nits for some HLG material (1,000 nits is an established normalization point for HLG to PQ conversions). Many non-HDR consumer TVs scale SDR to approximately the same peak white brightness level as the HDR anchor point (reference white at 203 nits).

For right now, most HDR productions use SDR graphics which mean they don't have to change their workflow. The SDR-to-HDR conversion LUT handles the mapping of graphics into HDR. By establishing an anchor point, HDR graphics producers could build elements and know where to map them -- when looking at a scope, they can reference those known white levels. This reference white level is extremely important to making the production consistent, especially for live material. Again, the overriding goal is to ensure that the colors and light levels delivered are consistent and aren't so uncomfortably bright as to be garish.

Figures 10 and 11 illustrate the concept of HDR shading, with no highlights (Figure 10) and highlights (Figure 11.) Highlights (like the sky in Figure 11) are typically a small part of the content. In Figure 10, an average scene from a game correlates on the scope image to roughly a 75% level, which indicates that the levels are set well. In Figure 11, a segment of the sky is in the shot, which immediately skews the corresponding scope level above 75%. Some of the scene reaches 100%, which is equivalent to 1,000 nits. The CIE1931 (an earlier version of the chromaticity measure for color gamut devised by the Commission Internationale de L'Eclairage) in Figure 10+11 maps the visible chromaticity range and shows how the current image occupies a specific color space. It is the colorful plot in the bottom middle of the scope screen capture. The plot for BT.2020 WCG imagery shows the picture information well inside the BT.709 triangle. The circle on both figures is a vector scope, which helps identify standard color targets which color bars should fall into.

HLG images are normalized to 1,000 nits as per ITU-R Recommendations during conversion to PQ so that when the signal is distributed, displays of all brightness levels will match the original intent and not be uncomfortable. Using this normalization level means that PQ's absolute mapping will preserve brightness levels at their originally-created levels seen by the production shader.

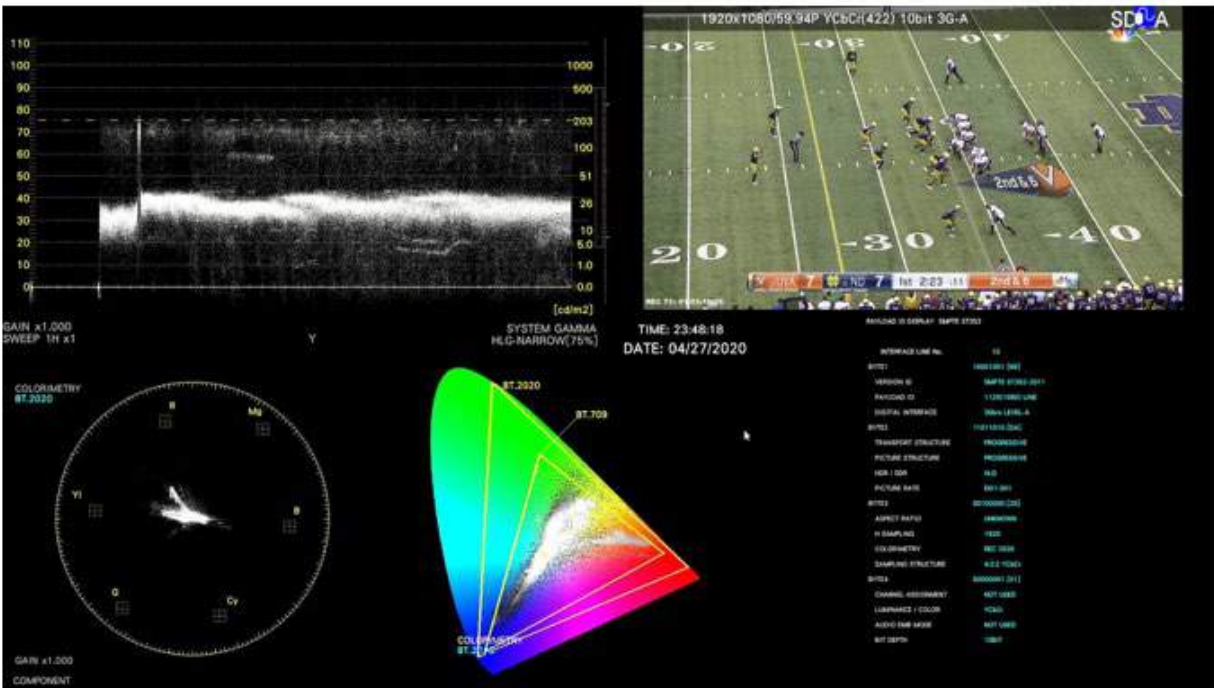


Figure 11 -: Examples of Shading: No Highlights




Figure 12 - Examples of Shading: Highlights

4. HDR/SDR Single-Stream Distribution

Figure 14 illustrates a simplified, single-stream HDR workflow for distribution. At the network origination point, all the production sources – SDR sources are de-interlaced, tone-mapped, scaled and color space converted. Since NBCU is agnostic to the HDR production format, native HDR formats (HLG or camera logs like S-log3) are converted to the final distribution format, PQ and resolution-scaled if necessary. Any third-party, native HDR content is also readied for transmission.

The process of cross-conversion from HLG to PQ has been defined in ITU-R Report BT.2390 (Figure 12) and produces an mathematically transparent conversion. It specifies a normalized peak white at 1,000 nit using an HLG system gamma of 1.2. Since HLG has a smaller color volume compared to PQ, we typically allow overshoots up to 109% for HLG (9% above peak-white) during conversion.

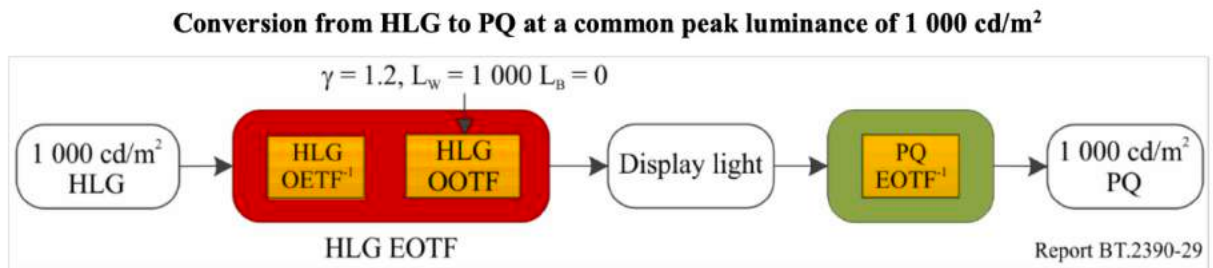


Figure 13 - This conversion produces a PQ image which is visually identical to HLG

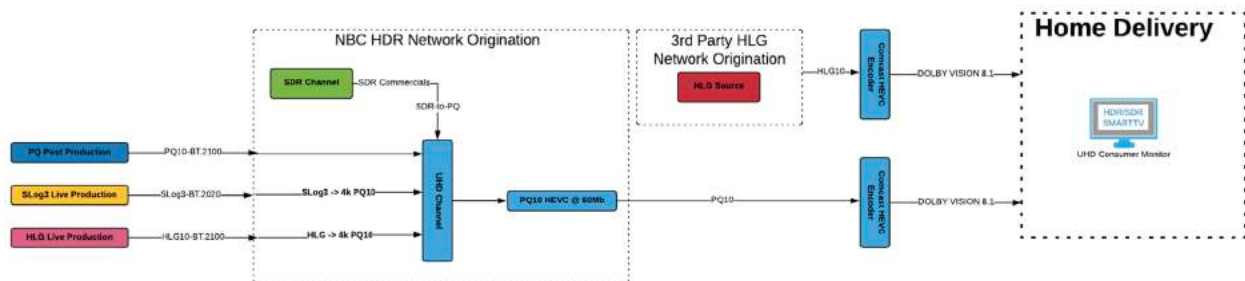


Figure 14 - A simplified view of the NBCU/Comcast production through distribution in PQ

The final distribution includes:

- Event production video: Converted from HLG/Slog3 to PQ
- Commercials: Converted from SDR to PQ
- Audio is encoded externally as Dolby ATMOS immersive audio and then embedded in the video prior to the HEVC encoder.
- Video is encoded using high efficiency video coding (HEVC) at 60-80 Mbps
- Mezzanine encode is transmitted to distribution partners (cable, satellite, and IP/OTT.).

NBCUniversal and Comcast use PQ for linear distribution because, for starters, it is the most common HDR system in-market: it's in every HDR display TV, and supported in OTT streamers like AppleTV and Roku; and, OTT providers like Netflix use it. It's used in PQ10, HDR10, HDR10+ (in each case, the "10" represents 10 bits, not the 8-bit transmissions of SDR). It can use static or dynamic color volume mapping (as mentioned earlier), which improves tone-mapping to displays with different brightness capabilities. Its conversions from other HDR formats are transparent, and normalized conversion from SDR to PQ preserves the original artistic intent.

Comcast is able to take NBCU's higher contribution rate HEVC video and re-encode it to a lower distribution rate while also adding Dolby Vision dynamic metadata to the HEVC stream. The Dolby Vision metadata provides scene or frame-based dynamic tone-mapping which provides benefits when the consumer display has different capabilities compared to the source. As televisions peak brightness

capabilities improve, the content producers may choose to improve the content as well but initially not all of the more affordable consumer displays will support the superior capabilities. Using dynamic metadata can improve tone-mapping from content to display which much of the consumer ecosystem can benefit from.

An additional benefit of mapping an HLG production to PQ for distribution is that we can take advantage of PQ's support of static or dynamic metadata for improved tone-mapping on less expensive displays that don't support luminance at or above 1,000nits. Finally, many HLG consumer displays clip above nominal video levels (100 IRE), but we can preserve levels above HLG-100IRE during the conversion because PQ supports up to 10,000nits(10x that of HLG). The IRE unit is used in the measurement of composite video signals.

In live-linear broadcast distribution, PQ, with proper shading, means that as displays get brighter, the images will remain at the original brightness unless the consumer manually changes settings. The original artistic intent is preserved and does not suffer from some other perceptual effects referenced in section 2.1 (Hunt, Stevens, etc.). Since some cinema or pre-produced content are currently shaded at levels up to 4,000 nits, relative systems like HLG, could be restrictive as displays improve. In short, for live-linear distribution, PQ's absolute mapping, combined with the productions white level anchor point, make beautiful imagery and preserves artistic intent all the way from the production to the consumers HDR display.

5. Conclusion

Producing and delivering HDR content for live linear workflows, like sports, is considerably more complex than it is for most file-based, VOD-styled approaches, which have post production resources that simply don't exist for live, on-the-fly material. Live productions must handle graphics, color correction, and luminance normalization in real-time, across what is usually a mix of SDR (which is essentially contemporary HDTV) and HDR cameras.

This paper details differences in system colorimetry which exist within the primary HDR systems, and specifically HLG) and PQ. It describes functional workflows for live delivery, including what are multiple levels of conversion for various signals throughout the infrastructure -- like graphics, video playback, SDR and native HDR cameras.

It has covered how format conversion and perceptual measurement work for HDR and SDR productions, how HDR production workflows are orchestrated for live content, and how those live linear workflows are distributed in PQ. Most contemporary broadcasters and cable providers are following a similar, if not identical, path: Producing in HLG or other camera-logs and distributing in PQ.

This paper also explained what it takes to build a functional, single-stream HDR-SDR workflow, at each stage of the pipeline, without compromising the artistic intent or quality of the distribution paths for either the HDR or the primary SDR feeds. In-depth research of color conversions were discussed and explored, as well as production layouts.

Finally, distribution requirements, delivery infrastructure and technical specifications were described and explored. In particular, we described why production environments tend to prefer to use the HLG transfer function, while distribution functions favor PQ.

Abbreviations

BBC	British Broadcasting Corporation
CIE	Commission Internationale de L'Eclairage
CRT	cathode ray tube
DLC	display light conversion
EOTF	electrical-optical transfer function
FIFA	Federation Internationale de Football
HDR	high dynamic range
HEVC	high efficiency video coding
HLG	Hybrid-Log Gamma
HVS	human visual system
IP	Internet protocol
IRE	Institute of Radio Engineers
JND	just noticeable difference
LUT	look-up table
NHK	Nippon Hoso Kyokai (Japanese Broadcasting Association)
OETF	optical-electrical transfer function
OTT	over-the-top
PQ	perceptual quantization
RGB	red, blue, green
SDI	serial digital interface
SDR	Standard dynamic range
WCG	wide color gamut

Bibliography & References

ITU: ITU-R Recommendation BT.709-6 (2015), “Parameter values for the HDTV standards for production and international programme exchange ,” International Telecommunications Union, Geneva

ITU: ITU-R Recommendation BT.2020-2 (2015), “Parameter values for ultra-high definition television systems for production and international programme exchange,” International Telecommunications Union, Geneva

ITU: ITU-R Recommendation BT.2100-2 (2018), “Image parameter values for high dynamic range television for use in production and international programme exchange,” International Telecommunications Union, Geneva

ITU: ITU-R Recommendation BT.2124-0 (2019), “Objective metric for the assessment of the potential visibility of colour differences in television,” International Telecommunications Union, Geneva

ITU: ITU-R Report BT.2390-8 (2020), “High dynamic range television for production and international programme exchange,” International Telecommunications Union, Geneva

ITU: ITU-R Report BT.2408-3 (2019), “Operational practices in HDR television production,” International Telecommunications Union, Geneva

ITU: H-series Recommendations–Supplement19, “Usage of video signal type code points,” International Telecommunications Union, Geneva.

Min, H., Kim, T., & Weyrich, JK. (2009). Modeling human color perception under extended luminance levels. *Association for Computing Machinery Transactions on Graphics, Vol. 28, No. 3*.
<http://reality.cs.ucl.ac.uk/projects/xlrcam/kim09xlrcam.pdf>

vooya video sequence player and YccIccSee color plugin. <https://www.offminor.de/>

Reducing The Tradeoff Between Performance And Management Using Container And Cloud-Native Approaches

A Technical Paper prepared for SCTE•ISBE by

Karthik Krishna

Senior Solutions Manager
VMware
Palo Alto, California
kakrishna@vmware.com

Jambi Ganbar

Director, Telco Partner Solutions
VMware
Berlin, Germany
jganbar@vmware.com

Table of Contents

Title	Page Number
Introduction.....	3
1. Business and Technology Drivers	3
1.1. Demand for Services	3
1.2. Edge Computing/IoT Opportunities	3
1.3. Technology Evolves, Virtualizes.....	4
2. The Virtualization Story	4
2.1. IT Legacy	4
2.2. DPDK and Fast Packet Processing.....	5
2.3. Architectural Considerations	5
3. MSO Options	5
3.1. Legacy Appliance-based.....	5
3.2. Virtualization – Evolved on Bare Metal	6
3.3. Native on Hypervisor	6
4. Next Gen Cloud	6
4.1. Next-Gen Architecture	6
4.2. Next-Gen Benefits	7
4.3. Next-Gen Capabilities	7
5. Data Plane Acceleration.....	8
5.1. Performance and Tradeoffs	8
5.2. Two Data-Plane Acceleration Options.....	9
5.2.1. Acceleration with Performance NIC	9
5.2.2. Performance NIC Test Results	10
5.2.3. SmartNIC.....	11
6. Next-Gen Cloud Reference Architecture.....	11
7. vCMTS Use Case	12
Conclusion.....	13
Abbreviations.....	13
Bibliography & References	15

List of Figures

Title	Page Number
Figure 1 – Next-Gen Cloud Architecture.....	6
Figure 2 – Performance NIC Architecture.....	9
Figure 3 – SmartNIC Architecture	9
Figure 4 – Acceleration with Performance NIC	10
Figure 5 – Acceleration with SmartNIC.....	11
Figure 6 – Next-Gen Cloud Reference Architecture.....	12
Figure 7 – Scale and Convergence in Next-Gen Cable Cloud	13

Introduction

Virtualization delivers efficiencies and creates new capabilities, while expanding the network cloud and blurring boundaries at the network edge. In the initial wave of virtualized functions and appliances, demonstrating that software-based infrastructure could be as robust as its hardware equivalent has been a paramount concern. Recent developments in data plane acceleration technologies, including SmartNICs, are playing an important role in meeting and exceeding performance requirements. Going forward, as services expand to meet opportunities, especially in the evolving network edge, the challenge will be to narrow the gap between performance and manageability. A Next-Gen Cloud model promises to minimize those trade-offs.

1. Business and Technology Drivers

1.1. Demand for Services

For change or evolution in a network to occur, there needs to be a business reason. A prerequisite for any change is demand for services. There also needs to be technology that can make the delivery of those services more valuable and profitable. But in regard to revenue generation, the cable industry is in a relatively strong position.

Demand for high-speed data remains strong. Since March 2020, subscribers have especially taken advantage of the upstream capabilities of DOCSIS networks. The industry's flexible content delivery networks (CDNs) are an attractive platform for businesses seeking to deliver automated and targeted advertising. Cable's telephony business has declined, as have all landline offerings, but the industry's MVNO business is up. There is also a growing opportunity to work with mobile operators on backhaul, mid-haul and fronthaul (Xhaul), as smaller cells proliferate in the 5G/LTE-A transition. MSOs continue to build their business services portfolio, both in the SMB market and beyond.

1.2. Edge Computing/IoT Opportunities

Among those existing services, perhaps the work being done in CDNs (manifest manipulation, formatting, encryption, etc.) is the leading example of computing at the network edge today. But at the intersection of wireless, consumer and B2B, new applications involving mobile and IoT devices that need similar services are emerging across numerous verticals. A few examples include:

- Cloud gaming – A billion-dollar industry that has grown during the pandemic, cloud gaming benefits from low latency associated with edge-computing; in effect, it “needs an edge.”[1]
- Surveillance, mapping – Distributed network sensors from security cameras, drones and IoT devices drive massive volumes of traffic across constrained wireless networks; one opportunity is for more integrated and AI-driven management.[2]
- Health care – From medical-grade wearables to RFID tags in hospital inventories, IoT is playing a key role in today's increasing digital healthcare system; it needs secure, robust connectivity.
- Smart cities – Municipalities aiming to deploy outdoor Industrial (I)IoT systems that connect remote sensors to network servers could leverage vast HFC networks.[3] From automated traffic lights to self-driving vehicles, what's needed is very low latency and very high reliability.
- Augmented reality – Related to gaming, AR applications are also playing a role in business and industrial settings; like other examples, it could benefit from low-latency and AI.

Cable executives see potential in several of these and other categories. But according to a survey conducted by Broadband Success Partners, they face obstacles, including a lack of operational support for

monetization and the existing network itself. When asked what is most needed to make edge computing happen, they have one leading answer: “Network infrastructure (physical and virtual) that is programmable and provides network performance information.”[4]

1.3. Technology Evolves, Virtualizes

The industry’s technology leaders have adeptly exploited the capabilities of the HFC plant with iterative versions of DOCSIS, while anticipating areas of growth, such as the need to meet demands on the upstream or low-latency requirements for Xhaul and other emerging services.[5][6]

Some MSOs have also begun to adopt cloud computing and virtualization, seizing opportunities for efficiency, innovation, lower OpEx, automation, improved management, elasticity and scale. The use cases involve program guides/UIs, CPE devices, CDN and most notably the CMTS/CCAP. In new cases following the distributed access architecture (DAA), the virtualized (v)CMTS is not only disaggregated, but its core software is loaded onto x86 servers and compute power deployed on remote PHY devices (RPDs) in the last mile.

The DAA initiative was driven less by new market opportunities than the need to manage costs, especially those involving power and headend/hub real estate. The goal has been to grow without launching a massive rebuild or upgrade cycle; in other words, to keep cost per bits stable while meeting escalating demand and maintaining stringent service performance requirements.

At the same time, the vCMTS has opportunities at the network edge. Here, again, the industry is being proactive.[7] One relevant trend is the move away from scale-out homogenous servers using x86 CPUs only, to scale-out heterogenous services that include x86 CPUs, as well as graphics processing units (GPUs), field-programmable gate arrays (FPGAs), and other elements.[8]

As usual, prospects arrive with additional challenges. Delivering services over virtualized and software-driven infrastructure within tight performance and cost parameters is no small feat. But there are other requirements to consider, such as agility, network scalability, security, operational support, product lifecycle and overall management capability.

2. The Virtualization Story

2.1. IT Legacy

Network virtualization is far from a new technology. The idea of a virtual machine (VM) goes back at least to the mid-1970s.[9] Two decades later, in the tech boom of the late 1990s, the time was right to focus on software that acted like a real computer with an OS, separate from underlying hardware. The first product of VMware, founded in that era, was a hosted hypervisor, which enabled users to set up and maintain VMs on a single physical machine.

This market primarily focused on IT, on business-critical applications such as data base systems and Web applications. The emergence of the Network Functions Virtualization (NFV) initiative in 2012 shifted attention to areas such as routing, firewalls and load balancing, all workloads requiring much higher packet rates and much lower packet loss rates. In the case of VMware, these requirements led to a “re-architecture” of the entire networking stack, from the virtualized network interface controller (vNIC) emulation to virtual switching in the device driver.[10] As it happens, there were ongoing industry initiatives collaboration to support such efforts.

2.2. DPDK and Fast Packet Processing

The Data Plane Development Kit (DPDK), founded in 2010 by Intel, became an important resource. Made available under an open source license, DPDK was created to be a “vendor-neutral software platform for enabling fast packet processing, upon which users can build and run data plane applications.”[11] The open source community at DPDK.org launched in 2013, and the initiative went under Linux Foundation management in 2017. Alternatives to DPDK are available; according to software development community StackShare, they include Beats, Riemann, LibreNMS, PRTG and Nagios XI .

For its part, the DPDK and its resources, which include libraries and drivers to implement network functions in commodity servers with commodity NICs, have driven industry-wide data-plane efficiencies. Partly inspired by DPDK techniques, VMware built and released an initial enhanced networking stack (ENS) that delivered a packet forwarding rate, in 64-byte packets, with 4 times greater efficiency, while maintaining a packet loss rate of less than 0.001 percent. Since then, performance has continued to improve.

2.3. Architectural Considerations

Other developments impacted the landscape. At the OS-level, virtualization enabled the delivery of resource-efficient software in packages called containers, which are associated with a release by Docker in 2013. Soon thereafter came an open-source container orchestration system known as Kubernetes, designed by Google and now maintained by the Cloud Native Computing Foundation (CNCF). While often pitted against each other, these approaches can all coexist. The latest VMware platform, for instance, allows using any combination of VMs and containers.

At the system level, disaggregated functions mean that control, data and management planes can be deployed across a distributed topology. Where to place what, however, follows basic rules. Because of inherent low latency and advances in processing, edge clouds offer performance advantages for data plane-intensive workloads; while control and management plane components (to the extent that they figure within a virtualized infrastructure) can be centralized with regional and global scope.

One example is distributing compute and virtualization via Flexible MAC Architecture (FMA), a CableLabs initiative. FMA provides the flexibility to place the MAC or compute portion of the CCAP anywhere from the DataCenter to Headend/Hub or in the outside cable plant. The FMA enables placement of the MAC of such a data-plane intensive workload closer to the Edge to address low-latency applications like AR/VR, gaming, or autonomous vehicles and centralize other functions like control and management plane. Such a flexible edge computing model enables smooth integration with wireless.

3. MSO Options

3.1. Legacy Appliance-based

The cable industry has several options regarding virtualization. Whether involving routers, CDNs, firewalls, CMTS/CCAP or other equipment, MSOs can expand their infrastructure as needed via standard appliance-based solutions. This largely status-quo approach, even with disaggregation, poses continued challenges to scaling and incurs operational costs and complexities.

3.2. Virtualization – Evolved on Bare Metal

Virtualizing on bare metal, i.e. on a single-tenant physical server, can deliver certain benefits, including performance and resource utilization; although challenges remain across a number of areas, including scalability, data persistence, networking, security, and management/operations.

3.3. Native on Hypervisor

This second road to virtualization enables a unified, converged architecture, with the benefits of higher availability, densification, scale, multi-tenancy and greater manageability. This “Next-Gen Cloud” approach also enables greater agility. Both virtualization approaches entail breaking with legacy models and acquiring new areas of expertise.

4. Next Gen Cloud

To expand upon the third option above, the Next-Gen Cloud, let’s consider its platform architecture, benefits and values.

4.1. Next-Gen Architecture

The Next-Gen Cloud schematically rests upon a standard Network Functions Virtualization Infrastructure (NFVI) and Cloud-Native Network Functions (CNF) infrastructure. (See Figure 1.) Above that is the Virtualized Infrastructure Management (VIM) layer, which enables handling containers and related infrastructure “as a service” (IaaS, CaaS). The Virtualization layer provides resource abstraction for computing, networking, and storage and gives the same experience across VMs and containers.

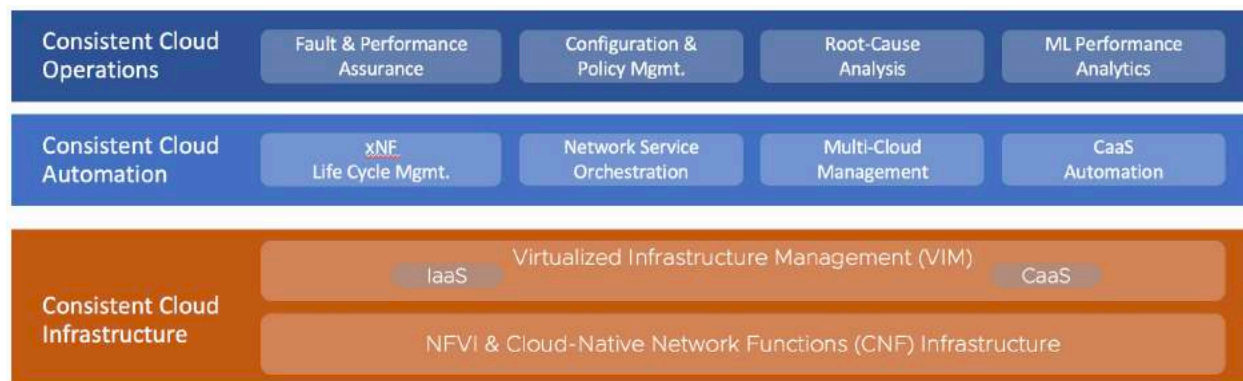


Figure 1 – Next-Gen Cloud Architecture

Cloud automation and operations reside above the infrastructure. Cloud Automation provides multi-layer automation for consistent operations across any network from the core to the edge to the access. Cloud automation enables smooth management and orchestration of the infrastructure layer and automates network functions' placement and lifecycle management over VM- and container-based infrastructure. Cloud Operations provides real-time assurance designed to simplify your network operations through holistic monitoring and performance management, providing comprehensive visibility and automation. Cloud operations in a virtual environment can integrate with existing NetOps tools and data, but they also enable cloud-native ways of fulfilling functions in the areas of fault and performance assurance, configuration and policy management, root-cause analysis (RCA), and machine-learning (ML) analytics.

None of these technologies work without revenue generation and/or positive ROI. Conceptually, at the top of the stack lie a wide range of monetization opportunities; and embedded within the stack are tremendous cost savings and management efficiencies.

4.2. Next-Gen Benefits

The benefits of this approach derive from the interplay of the platform's capabilities and performance in these areas:

Scalability – Orchestration enables elastic scaling of cable workloads, while consistent dimensioning and improved hardware utilization reduces space and power required in headend and hubs. The platform brings uniformity while maintaining security, networking and monitoring requirements.

Operational Efficiency – Having a common platform enables deployment and management of multiple Kubernetes clusters, along with VMs. Sophisticated software, which leads to fully automated services, compares favorably to an inefficient and complicated manual process of managing physical storage devices and using SNMP for monitoring to the NOC. It also provides a consistent experience with a simplified life-cycle management.

Faster Time to Market – This kind of platform allows MSOs to respond quickly to emerging edge compute and other opportunities through an extensive, multi-vendor partner ecosystem, as well as integrated service blueprinting, deployment, monitoring and management.

Workforce Efficiency – Remote deployment and high-performance, self-healing capabilities can reduce truck rolls and the need for on-site service technicians, whether in the access plant, hubs, headends or data centers.

Reduced TCO – Lower CapEx results from increased throughput and performance; reduced OpEx, from software-driven infrastructure and service assurance.

Safe and Easy Convergence – Consistent infrastructure, automation and operations enable easy convergence, with reduced risk of misconfiguration; hybrid VM and container-based applications feature advanced networking and security features.

4.3. Next-Gen Capabilities

Agility. With a hypervisor, a new bare-metal server can connect to the container domain in minutes. That agility is one of the reasons that major public cloud providers use hypervisors to run their container services. Such agility and scalability work because the hypervisor de-couples and abstracts software from the hardware. The Next-Gen Cloud delivers over-subscription capability, improving utilization of the underlying hardware. This compares favorably to the standard practice of racking and stacking, re-segmenting the network and then testing that the segmentation was correct.

Networking. The Next-Gen Cloud implements a single underlay network on VMs to provide end-to-end connectivity and management for both containers and traditional applications. A single underlay network makes it easier to connect containerized applications to traditional, non-containerized components like databases; simplifies network management with centralized policies and advanced security; and enables selecting the overlay network and the service mesh that works best for containerized applications.

Security. Containers and hardware virtualization not only can, but very frequently do coexist and actually enhance each other's capabilities. VMs provide strong isolation, OS automation and an ecosystem

of solutions. They enable secure and efficient running of containerized applications in production. Containers being run on hypervisors can take advantage of security innovations, including micro-segmentation, which enables security architects to apply security controls and deliver services at the individual workload level.

Manageability. A comprehensive, flexible platform allows you to deploy and manage multiple Kubernetes clusters as well as to manage, patch and upgrade the container host OS. All these capabilities empower you to run traditional and containerized workloads on a common infrastructure, while ensuring optimal performance and preventing interference between workloads.

Performance. Empowered by vSwitch and Workload Acceleration functions, the Next-Gen cloud's hypervisor can provide more efficient overall workload performance for containers than Linux systems running on physical hardware. The platform uses advanced scheduling algorithms that enable employing modern DPDK packet processing, allocating CPU cycles for efficient networking, and faster packet processing to optimize all workloads. The Next-Gen cloud also employs hardware offloading techniques to dedicate all the compute for workloads and can leverage SmartNIC for infrastructure computing.

5. Data Plane Acceleration

5.1. Performance and Tradeoffs

To look more closely into performance, let's first set the stage. Workloads in service provider solutions involve the control and signaling plane, the data plane, and the management plane. A network element (NE) traditionally was built to handle control and data plane functions in one box, with its own element management system (EMS). With the arrival of SDN, the data and control planes began to separate. Then with NFV, multiple functions within each NE were broken into multiple VNFs, providing service providers with choice, enabling best-of-breed solutions with more software-based control.

Of the three workloads, those on the data plane need to meet the most stringent requirements. (When the data plane is lost, the subscriber notices.) These demands led to the development of software accelerators that put applications needing fast packet processing in the 'fastpath' of a compute host. Management and low-priority applications could go via 'slowpath.' Yet the arrival of virtualization made it difficult to attain expected performance levels, because the hardware used was based upon commercial off-the-shelf (COTS) servers.

To be sure, virtualization correctly done continued to have significant benefits, primarily: those derived from the secure pooling of networking resources across multiple hardware units. A virtualization layer with a virtual distributed switch that spans multiple hosts (servers) can live-migrate workloads and efficiently distribute resources on demand. To gain performance, service providers used two techniques: single root I/O virtualization (SR-IOV), a concept introduced by Intel that allows for bypassing the virtualization layer when it relates to the data path; and DirectPath I/O, which allows for direct access to a physical NIC.

These techniques enabled faster packet processing but came at the expense of not being able to use a wide range of critical features, all provided by the best-in-class virtual infrastructure. Applications built for these pass-through mechanisms also faced security and cloud-ready challenges. Another technique that gained popularity was CPU pinning, which led to locked-in hardware, driving up CapEx and OpEx costs.

Among the tools enabled by the DPDK initiative was a software library that helped enhance performance by allowing for optimized packet allocation across DRAM channels. The principle being that allocation of memory from local nodes and cache-alignment of objects can lead to superior performance. The upshot

of these developments was an accelerated vSwitch, now a key to the Next-Gen Cloud model, which overcomes the drawbacks of SR-IO and DirectPath I/O and obviates the need for other techniques, such as CPU pinning.

5.2. Two Data-Plane Acceleration Options

With that background, we can now consider two data-plane acceleration approaches, both involving PCIe NIC devices. The first involves offloading with a standard Performance NIC; and the second, doing so with SmartNIC. Both cases use an x86 host server and the accelerated vSwitch, which supports two configuration modes: Standard, for use with any management or control-plane application; and Enhanced, for use with data-plane intensive applications. In the first approach, the vSwitch is running on the hypervisor, and in the second, it has been offloaded onto the SmartNIC and can leverage separate embedded NIC cores. (See Figure 2 and Figure 3.)




Figure 2 – Performance NIC Architecture




Figure 3 – SmartNIC Architecture

5.2.1. Acceleration with Performance NIC

Operators have several ways to boost performance, not only in the software realm. We will point out five areas. (See Figure 4.) Let's begin with decisions involving hardware elements:

- (1) NIC choice – To meet 10G requirements, first choose Performance NICs with TCP Segmentation Offload (TSO) and Checksum Offload (CSUM) capabilities, as well as overlay support.
- (2) CPU Choice – Devote a high number of cores for workloads and find the right balance between CPU speed, core count and wattage requirements.
- (3) Server architecture – Increased performance is a function of the NUMA balance, choice of the right PCIe and server architecture. Another key consideration is server immutability, in which servers are never modified after being deployed.

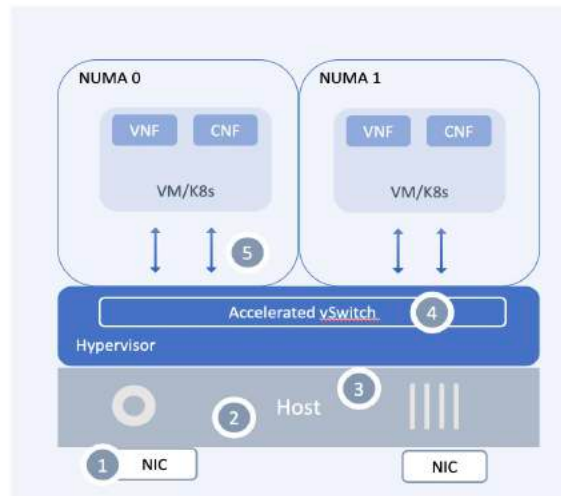


Figure 4 – Acceleration with Performance NIC

With the right mix of hardware, performance can be further enhanced through strategic choices in software implementation (enabled by Next-Gen cloud):

(4) vSwitch Acceleration – Based on Intel DPDK principles, an accelerated vSwitch provides very high throughput, low latency, low jitter, near zero-packet loss, characteristics that are essential to run data plane-intensive workloads, whether at the Core, the Edge or the Access. The accelerated vSwitch achieves higher packet performance by freeing the network to work on any server or modern NIC; dedicating CPU cycles for networking; and deploying faster switching with flow cache, lockless datapath and faster packet processing (SSE). The vSwitch forwards traffic between components running on the transport node (between VMs/VNFs or between containers) or between internal components and the physical network.

(5) Workload acceleration – Other options include using an accelerated vNIC, dedicating (pinned and isolated) vCPUs; topology awareness, and Huge Pages support.

Powered by the accelerated vSwitch and an accelerated workload with the right combination of NIC, CPU and server hardware, this approach can handle the various workloads in use today. VNFs in all types of form factors – whether VMs, containers, or micro services – can run on this kind of enhanced solution. Being agnostic to the application, the vSwitch provides operators with the ‘freedom of choice’ to pick any VNF that would meet the solution needs.

5.2.2. Performance NIC Test Results

Tests of Performance NICs, which can flag latency- and jitter-sensitive applications and choose selective vCPU pinning, generated results well within acceptable parameters. In a traffic stream starting and terminating with Spirent TestCenter software, with latency that included both transmission delay introduced in forward and reverse paths as well as software processing times, the average latency contributed by virtualization functions amounted to less than 30 microseconds; for jitter, less than 10 microseconds. Both are negligible amounts. When testing throughput of varying packet sizes, the actual throughput approximated the line rate, with a representative bundle of video-heavy packets generating a speed of 3.98 Mpps and 4.5 Mpps, respectively, for 3- and 4-core implementations.

5.2.3. SmartNIC

There is a growing consensus that CPU cores are such a precious commodity that they should never do network, storage, or hypervisor housekeeping work, but rather focus on core computation. That means network offloads and storage offloads need to be mainstream in the coming years, creating an even more asymmetric and heterogeneous processing environment than many are envisioning down the road. Throwing more and more compute at the problem is unsustainable. The goal is to iterate fast, but I/O processing is unable to keep pace with compute processing. SmartNICs enable offloading all or most of a virtual switching stack or a large chunk of a distributed storage stack.

A SmartNIC is a high-performance NIC, equipped with general-purpose compute cores capable of running an OS and general-purpose applications and workloads. It is both a NIC, with flow-match hardware capabilities, and a mini server. In architectural terms, its purpose is to move server management to SmartNIC cores and offload the entire vSwitch/hypervisor. As a result, the host x86 remains dedicated to workloads, increasing its efficiency. The host x86 can run either a VM/container or a bare-metal.

One of the big reasons to use a SmartNIC is not only to save CPU resources, but to scale in different ways, for instance by saving power and space. Offloading the heavy compute enables an operator to free up the CPU for other things, such as multi-access edge computing (MEC) in 5G. Another key to a full CPU offload is the potential to implement policy and priorities straight on the input, without having to use the OS at all, which makes for a powerful story.

Other benefits include a clear-cut between infrastructure and application, air-gap security, cost management, application acceleration (e.g. security, load balancers, etc.). In organizational terms, the SmartNIC clearly splits ownership and enables multi-tenancy. (See Figure 5.) SmartNICs are positioned to handle workloads that are scaling to meet 10G requirements.

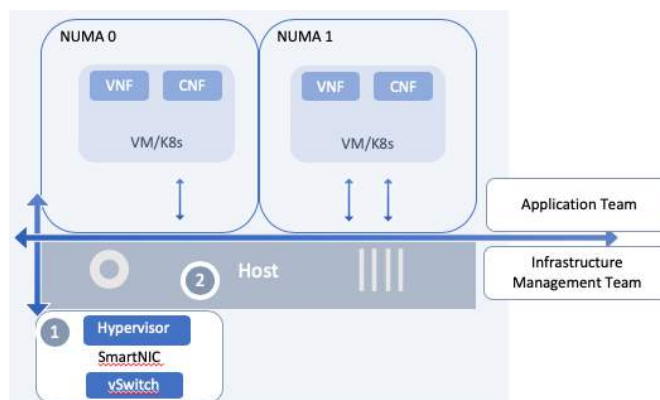


Figure 5 – Acceleration with SmartNIC

Because SmartNICs come equipped with onboard local persistent storage, a large amount of DDR RAM, multi-level caches, PCIe root complex, virtualized device functions and I/O capabilities, they have significant potential for use cases that extend beyond network acceleration and vSwitch offloads. Those cases could include but are not limited to management plane offload and virtual device acceleration.

6. Next-Gen Cloud Reference Architecture

The reference architecture for Next-Gen Cloud includes five tiers: from physical infrastructure to VM and container-based platforms to resource orchestration to cloud automation to solutions. (See Figure 6.) On

the solutions tier, the vCMTS is only one of many, which also include L2/L3 solutions, SD-WAN and various video-centric services. Being extensible to emerging solutions and business models is one of the most compelling attributes of the Next-Gen Cloud.

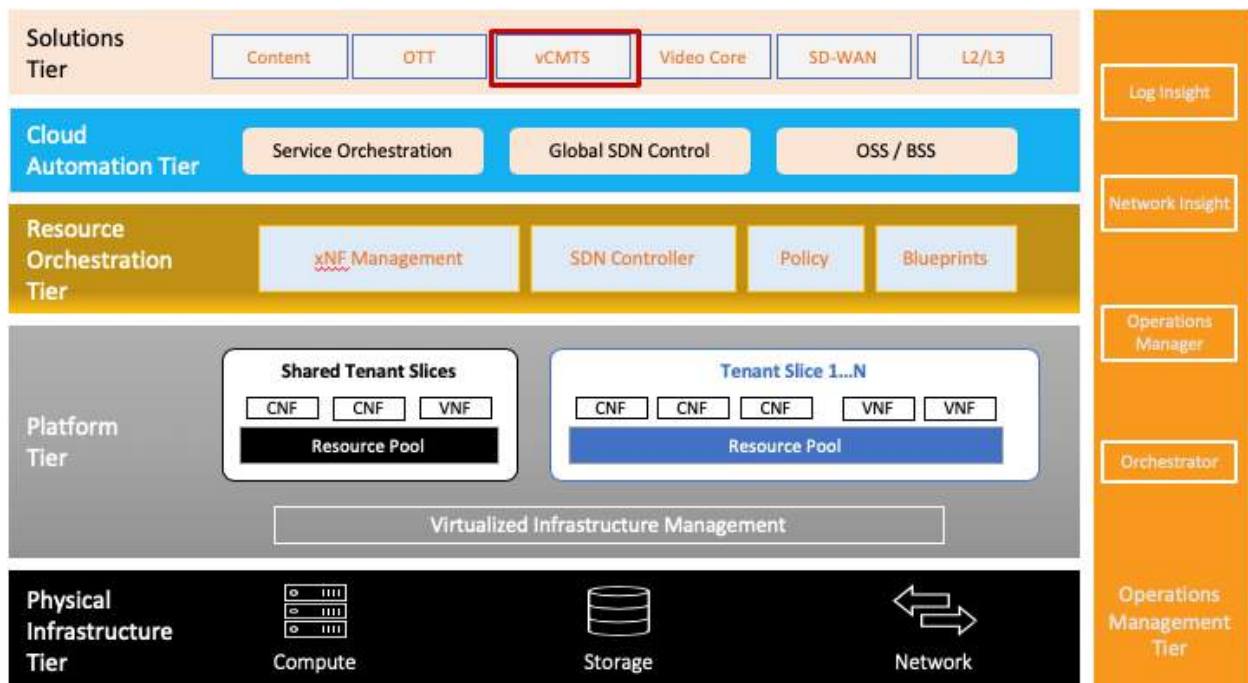


Figure 6 – Next-Gen Cloud Reference Architecture

Moving from that top tier down the stack, we note that it features centralized control and management, including embedded automation and optimization. Its highly flexible infrastructure as code (IAC) model is characterized by a heterogeneous runtime with Network Function and Resource Isolation. Finally, it rests upon commodity hardware and storage and leverages vertically integrated

7. vCMTS Use Case

Compared to the mobile industry, where NFV has proved to be effective in 4G LTE production deployments and has become the basis for building 5G networks across extensive footprints, virtualization in the cable industry has taken a more gradual and piecemeal approach. The initial deployments of the vCMTS occurred after MSOs assessed different approaches to DAA and enhancements to DOCSIS. It is beyond our scope to assess those ongoing rollouts, but it is noteworthy that one summary practical lessons from a DAA deployment with a vCMTS last expressed concerns at the operational and management layer, in particular the need for tools.[12]

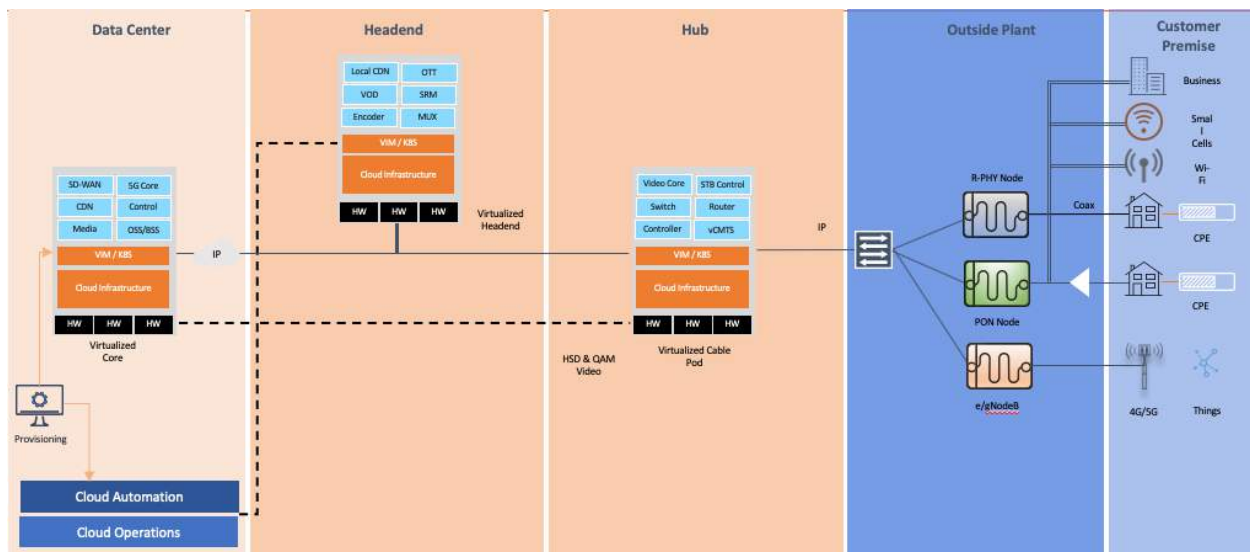


Figure 7 – Scale and Convergence in Next-Gen Cable Cloud

If manageability is an issue with the current vCMTS model, in which the R-PHY node is the only element in the outside plant, will it not become more so if and when MSOs deploy virtualized PON nodes and eNodeB devices? (See Figure 7.) Moreover, in addition to services delivered to residential CPE endpoints, those involving enterprises, small cells, Wi-Fi, 4G/5G and the proliferating number of edge compute business cases also need to be managed. Conducting efficient root-cause analysis over a network with no centralized control and integrated operations management is just one challenge we could mention.

On the other hand, with Next-Gen Cloud model, there is ongoing innovation in management, automation and OSS. Virtualization changes the way we can set up, handle faults, share resources, and more.

Conclusion

Growing demand for connectivity, compute and storage at the network edge is an exciting prospect for MSOs shifting to a distributed and virtualized network infrastructure model. Advanced data plane acceleration that can drive high performance over software-based infrastructure is also an encouraging development. Less exciting is the prospect of building more operational siloes to handle a growing menu of use cases and services, which defeats one of the fundamental reasons for the drive toward virtualization in the first place: “the principle that it is no longer necessary to solve for the platform and runtime layer below network applications in a different and particular way for each additional application – with all of the attracted cost, complexity, and operational management overhead that differentiation implies.”[13]

Abbreviations

4G LTE	fourth generation, long-term evolution
5G	fifth generation technology standard for cellular networks
AI	artificial intelligence
AR	augmented reality
CaaS	container-as-a-service

CCAP	converged cable access platform
CMTS	cable modem termination system
CNCF	Cloud Native Computing Foundation
COTS	commercial off-the-shelf
CPE	customer premises equipment
CPU	central processing unit
CSUM	checksum offload
DAA	distributed access architecture
DDR	double data rate
DOCSIS	data over cable service interface specification
DPDK	data plane development kit
DRAM	dynamic random access memory
eNodeB	E-UTRAN node B, or evolved node B
EMS	element management system
ENS	enhanced networking stack
GPU	graphics processing unit
HFC	hybrid/fiber coax
I/O	input/output
IAC	infrastructure as code
IIoT	industrial internet of thing
IoT	internet of things
LTE/A	long-term evolution/advanced
ML	machine learning
MEC	multi-access edge computing
MSO	multiple system operator
MVNO	mobile virtual network operator
NE	network element
NFV	network functions virtualization
NFVI	network functions virtualization infrastructure
NIC	network interface controller (or card)
NUMA	non-uniform memory access
RAM	random access memory
OS	operating system
PCIe	peripheral component interconnect express
RCA	root-cause analysis
RFID	radio frequency ID
RPD	remote PHY device
SD-WAN	software-defined wide area network
SP	service provider
SR-IOV	single root I/O virtualization
SSE	streaming SIMD (single instruction multiple data) extensions
TCO	total cost of ownership
TSO	TCP segmentation offload
UI	user interface
vCMTS	virtual CMTS
VIM	virtualized infrastructure manager
VM	virtual machine
vNIC	virtual NIC

Bibliography & References

- [1] Alan Evans, “Why Gaming Needs an Edge,” SCTE-ISBE, 2019.
- [2] Sandeep P. Chinchali, et al., “Neural Networks Meet Physical Networks: Distributed Inference Between Edge Devices and the Cloud,” HotNets-XVII, November 15-16, Association for Computing Machinery, 2018.
- [3] Charles Chapman, “How HFC-based Industrial IoT Gateways Improve Performance of Remote IoT Sensors,” SCTE-ISBE, Cable-Tec Expo, 2019.
- [4] David Strauss, “Cable and edge computing: The rest of the story,” Light Reading, May 18, 2020.
- [5] Jennifer Andreoli-Fang, John T. Chapman, et al., “Blueprint for Mobile Xhaul over DOCSIS,” SCTE-ISBE, 2019
- [6] Greg White, et al., “Low Latency DOCSIS: Overview and Performance Characteristics,” SCTE-ISBE, 2019
- [7] Omkar Dharmadhikari, “Moving Beyond Cloud Computing to Edge Computing,” CableLabs, May 1, 2019
- [8] Randy Levensalor, “Give Your Edge and Adrenaline Boost: Using Kubernetes to Orchestrate FPGAs and GPU,” CableLabs, Jan 28, 2020
- [9] Gerald J. Popek, Robert P. Goldberg, “Formal Requirements for virtualizable third generation architectures,” Communications of the ACM, July 1974
- [10] Moiz Alam, et al., “High Packet Rate Networking on vSphere,” VMware
- [11] DPDK Project Charter, dpdk.org/charter
- [12] Asaf Matatyaou, “Practical Lessons of a DAA Deployment with a Virtualized CMTS,” SCTE-ISBE, 2019
- [13] Andrew Bender, “A Roadmap for Virtualization in HFC Networks,” SCTE-ISBE 2019

Building a Business Service in the Cloud

A Technical Paper prepared for SCTE•ISBE by

Adrian Beaudin
Senior Architect
Akamai
Cambridge, MA
+1 613 670 8451
abeaudin@akamai.com

Bruce Van Nice
Senior Product Marketing Manager
Akamai
Santa Clara, CA
+1 650 381 6074
hvannice@akamai

Table of Contents

Title	Page Number
1. Introduction.....	3
2. On the Path to Cloud.....	3
3. Service Considerations	4
4. Data Plane versus Control Plane.....	4
5. Subscriber Experience	4
6. Security and Privacy	5
7. Integration into Provider Systems.....	5
8. Cloud Considerations.....	6
9. APIs.....	6
10. Customer Support and Operations	7
11. Summary	7

1. Introduction

Saturation of traditional markets is forcing ISPs to evaluate new strategies to drive revenues while continuing to deliver a superb user experience and improve returns on capital. Differentiating service offerings is critical to achieving these goals. Commodity connectivity doesn't always motivate subscribers to stay with a service or offer any incentive to pay more for it. Subscriber-facing services that enhance loyalty and increase "stickiness" can help.

Providing managed services and moving up the stack can find a receptive audience amongst business customers. Properly targeted value-add subscriber services can increase ARPU and differentiation and improve retention. But mitigating the costs of this specialty care will drive the use of more efficient, deployment solutions. One of the challenges delivering subscriber services is controlling costs to maintain margins in a world where specialized staff to oversee service deployment and operation are costly and difficult to find. Cloud services deserve consideration to reconcile these conflicting objectives

Cloud services can help contain costs and improve service agility (time to market). They've become mainstream for enterprises, but ISP service enablement in the cloud is a different proposition than typical enterprise IT applications. This paper will offer perspectives from the developer of a cloud service that enables ISPs to deliver security protections for businesses.

2. On the Path to Cloud

Starting in early 2013 a team at Nominum (acquired by Akamai in 2017) started down the path of building a cloud-based security service offering for ISPs. The security service consisted of DNS resolvers and a platform that supported functions such as gathering data, distributing policies, and supporting a portal. As with most development projects there was an evolution with several early, admittedly highly tactical, initial steps.

The first effort was to deploy the software components using the services of a public cloud provider. It was activated at a security conference to provide production DNS resolution (rDNS) services on an open network. Integration of dynamic threat intelligence identified devices with bot infections and there was a portal available for attendees to check their status.

Learnings from this pre-proof of concept led to deployments to support customer pilots of the security service. At the time, getting pilots running at ISPs involved lengthy engagements to get equipment and technician/administrator time in labs. Running these early solutions in the public cloud bypassed this time and materials bottleneck and permitted a faster time to market metric.

Some of the pilots in the public cloud were opened up to run limited production network traffic, so the next logical step was to run more scaled production workloads. The first foray was to support a customer who needed a network protection service immediately but had to contend with a 90 day window to procure and provision hardware in their network to support it. Successful execution of this service in the cloud provided validation and the process of formalizing the service began. All of the early experiments yielded useful insights that fed specifications for development teams ongoing work. Migrating functions to the cloud not only yielded efficiency from the network provider, it also removed the need for the IT team to procure, install and manage the servers and apps attendant to those functions.

3. Service Considerations

The first step for the product and development teams was agreeing on the parameters of the “service”. The following sections highlight the primary service considerations.

4. Data Plane versus Control Plane

Security is about examining network traffic to determine whether it is malicious or legitimate.

Approaches that operate in the data plane implement inline packet inspection to evaluate traffic. In the past it was possible to get visibility into most of the traffic on a network, but the predominance of encryption has introduced significant limitations. Inspecting packet traffic at line speeds has always been a costly operation, and inspecting encrypted traffic adds even more costs and operational overhead.

DNS filtering is an alternative that operates in the control plane. Incoming subscriber queries can be matched against dynamic threat intelligence provisioned in resolvers, and policies can be applied to manage unwanted traffic. The team recognized attractive scaling capabilities since there is no need for pervasive filtering of network traffic (and decryption) and it can be layered on infrastructure already deployed and managed. Threat coverage can be expanded by selectively forwarding suspicious traffic (typically domain names that point to both malicious and legitimate web resources) to proxies for further inspection. Experience has shown this is a small percentage of traffic, usually around 2%.

An obvious question given the emergence of DNS encryption standards is whether they make DNS traffic opaque to network operators. The new standards, DNS over TLS (DoT) and DNS over HTTPS (DoH) define encrypted transport for queries between stub resolvers implemented in client devices, and resolvers deployed by network operators (ISPs, MNOs, enterprises, Wi-Fi, etc.). Operators of resolvers using encrypted transports still see queries in the clear and services provided by the resolver function as they would with unencrypted transport.

5. Subscriber Experience

A subscriber’s Internet experience is closely tied to latency. In this case DNS resolution is an important part of the solution and the network architecture called for resolvers situated at the network edge as close to subscribers as possible to minimize transit delay. They were dimensioned to align expected subscriber densities with resolver performance (queries per second) to rightsized capacity.

Experience with resolution infrastructure at more than 100 large ISPs worldwide for nearly 20 years showed in today’s fixed networks (2020) a subscriber account can generate 10,000 queries per day. Growth has averaged about 20% per year. This provided metrics for sizing subscriber demands on resolvers, and along with a best practice targeting QPS performance at 50% CPU capacity to minimize latency under load, and 5X headroom for growth yielded necessary resolver capacity.

Services that can be customized have more potential to actively engage customers and provide an incentive to remain loyal since they’ve made an investment in configuring the service to meet their unique requirements. A natural extension of filtering to prevent malicious activity is customer defined filters to block unwanted content for families (parental controls) or businesses (Acceptable Use Policies).

A portal user interface is the subscribers window into their household security posture and another contributor to satisfaction with the service and thus value. Key words for the UI team were “rich, relevant, simple, comprehensible”. Menus define web filters and integrated device management capabilities

simplify configuration of user/device specific profiles through device discovery, registration, and pairing functions. Reports display Internet usage and security threats deterred. Scaling, especially with personalization, introduced complexity in the portal infrastructure. Underlying cloud services had to support secure access to millions of unique portal instances, as well as distribution of policy/preferences in near real-time to instantiate subscriber preferences, and collection of data to populate displays for each subscriber.

6. Security and Privacy

The team recognized concerns about security have been an inhibitor for adoption of cloud services and designed layers of defenses to protect the different components of the service. Portals for both administrative functions controlled by the ISP (discussed below), and subscriber-facing functions are protected with each providers Identity and Access Management system. A Web Application Firewall (WAF) protects server and user side APIs (discussed below). Perimeter defenses add another layer of protection.

Privacy is a dominant issue virtually everywhere in the world today and especially in developed countries. Privacy regulations in many parts of the world such as the EU General Data Protection Regulations (GDPR) define frameworks for the permissible collection and management of personal data. Security services and similar services like content filtering are subject to privacy regulations and the development team established several overriding principles followed in the design of the system:

- Only data that is necessary to operate the service, and provide reports covering its operation to subscribers, is processed. Data is not retained longer than necessary for any purpose.
- Data is pseudonymized where possible (e.g. subscriber IPs) in ways that render it impossible/impractical for a 3rd party to compromise subscriber privacy.
- Subscribers can guide processing of their personal data by expressing their preferences in an online portal. They can see the results of filtering (processing) at any time.

Additional data governance considerations motivated the definition of data processing and retention policies. Software built into the components of the service encrypts data in motion. AWS Elastic Block Storage (EBS) encrypts data at rest so the whole partition is encrypted. Further, data in many cases had to be stored and processed in the country where it was collected (data residency) so in-region cloud services are used. Providers access to all of their data at all times on a real time basis was built into the system. A function was also added that allows providers to zero out their data on demand.

End user requests about their data are currently handled by customer Service Representatives (CSRs) and forwarded to support channels. Data is also available in the system for provider operations teams to integrate with in-house systems to automate this function.

7. Integration into Provider Systems

Cloud services simplify deployment of ISP services by off-loading functions that providers would otherwise have to install and maintain in their networks. Integration with provider systems is still required to:

- Provision Subscriber identifiers, typically within an abstract subscriber ID that's not tied to an identifier that might change (like an IP address or MAC address) and that does not reveal Personally Identifiable Information (PII)

- Make policy changes requested by subscribers through a customer service representative (CSR) or self-service by the subscriber too
- Manage IP address updates in real time through RADIUS messages or DHCP logs. This is necessary due to the decision to allow subscribers to configure their unique content filtering preferences. Their policies have to follow them through address changes.
- Support SSO to allow for integration with subscriber and applications portals for CSR (these portals are described below)
- Connect to data collection and management systems. There's considerable built in flexibility to stream protocol and service logs to a central repository to support internal data collection and reporting functions for:
 - Customer support
 - Service adoption
 - Service utilization
- Enable functions required by BSS such as reporting on service adoption

8. Cloud Considerations

In early 2017 the initial cloud release was ready for production networks. At a high level there were 4 major components to the service to support in the cloud: DNS resolution, data transport, policy management, and subscriber and operator interfaces. Teams first explored how to scale the service in the cloud for daily usage patterns and peak events while accommodating growth in demand. Established metrics for the resolution component, discussed above, were extrapolated to build out the other components of the system.

Given the criticality of availability/reliability with respect to user experience a decision was made to deploy a single stack for each provider, versus taking advantage of potential gains from multi-tenancy (although these were never really measured so it's not clear they exist for this kind of service). Isolation has numerous positive implications, one customer can't take down the service, performance is more predictable, there's less data/privacy exposure and data can be repatriated. For resilience the service is deployed in fully redundant stacks in two different public cloud regions. A single region or service will not create an outage.

A Service Level Agreement tracks average availability on a monthly basis from telemetry data available through administrative service interfaces. Service credits are calculated when outages exceed an availability baseline.

The service was also designed so providers could deploy some components in their own networks in a hybrid configuration, or deploy the whole stack.

9. APIs

Early production deployments of the cloud service quickly highlighted features and functions that were necessary and not yet implemented. APIs rose to the top as essential. Provisioning interfaces were needed to integrate with provider operational processes and systems. Two provisioning APIs were created, one an abstraction and one low level. AWS APIs were a source of inspiration and insights for the development teams and used as the model for their efforts. Providers should hold all their cloud service vendors to a similarly high bar!

One of the provisioning APIs powered the subscriber portal so providers can build customized, purpose-built portals and match the user interface look and feel to other portals subscribers have access to. This API was beneficial from a vendor perspective as well since it reduced the burden developing differentiated features requested by individual customers.

The ability to provision subscribers/services was also exposed in this API - objects that represent subscribers and the services they have could be created. Portal authentication integrates with AAA such as RADIUS, LDAP can use SSO. A third API was created for IP to subscriber mapping using dhcp log scraping, radius accounting message ingestion or a REST API.

Several requirements were established for APIs

Proper documentation - Use swagger to generate docs in its standard format. Offer pdf docs with integration examples.

- Versioning. Enable critically important backwards compatibility for integrations.
- Simple API calls. Make heavy use of abstraction to minimize complexity.
- Read (GET) calls. Verify subscribers are properly provisioned in scaled, high throughput installations.

10. Customer Support and Operations

Customer support functions typically take advantage of APIs for subscriber provisioning and IP tracking, supplemented with logging information gathered by the service. A portal was designed to allow CSRs to look up subscribers, and their configuration (IP addresses, policies) as well as when they last generated traffic with the service. A pre-built portal was created, and APIs can be used to create a custom portal.

Another portal was created to provide global reports on the service such as Top infections, Top blocks, and service adoption. This reporting data can also be integrated with BSS or other custom interfaces used by a provider. Access to both portals is controlled with RBAC and can be integrated with an external directory such as LDAP.

11. Summary

Value-add subscriber services are becoming more strategic as ISPs look for new ways to grow revenues. Controlling costs for service enablement is essential, and reducing time to market is increasingly mandated, even as it's harder to find and retain specialized staff, and rigid processes to deploy new functions raise costs and inhibit agility.

Widespread availability of higher broadband speeds to both enterprises and residential customers has greatly facilitated the feasibility of control or data plane, cloud-based services. When typical customer locations have access to 100Mbps of bandwidth or greater, even with conventional latency the performance is high enough to allow remote and cloud compute as part of routine communications.

For these reasons cloud services deserve consideration to reconcile these conflicting objectives. They've become mainstream for enterprises, but ISP service enablement in the cloud is a different proposition than typical enterprise IT applications.

Product and development teams at Akamai created a cloud based security service designed to make it easy for ISPs to target customers with a new offer. Focus and 20 years' experience drove design considerations to:

- Enable the best possible subscriber experience, minimizing latency, providing a high degree of customization, and effective security with “just works” simplicity
- Build scale, availability and resilience into cloud-based components that support the service
- Ensure security and respect for privacy for the service itself, staff who operate it, and subscribers who use it
- Integrate with provider systems to automate provisioning, support monitoring, and collect data for business systems and to meet regulatory requirements

DNS Encryption: Exposure or Opportunity?

A Technical Paper prepared for SCTE•ISBE by

Mark Dokter

Senior Product Manager
Akamai
Toronto, Canada
+1 613 670 8451
mdokter@akamai.com

Bruce Van Nice

Senior Product Marketing Manager
Akamai
Santa Clara, CA
+1 650 381 6074
hvannice@akamai

Table of Contents

Title	Page Number
Table of Contents	2
1. Introduction	3
2. DNS Encryption Protocols.....	3
3. DNS Encryption and ISPs	4
4. DNS Encryption Client Implementations	5
5. Provider Impact of DNS Encryption Clients.....	7
6. Operational Impact of DNS Encryption	8
7. Summary Action Plan.....	10
Abbreviations.....	11
Appendix: DoT and DoH Client Implementations.....	12
Bibliography & References.....	14

List of Figures

Title	Page Number
Figure 1 – DNS over TLS defines encrypted transport between stub resolvers and resolvers using TLS	4
Figure 2 – DNS over HTTPS defines encrypted transport between stub resolvers and resolvers using HTTPS.....	4
Figure 3 – Encrypted transport termination with query processing in the clear.....	5
Figure 4 – Resolvers can be equipped with threat intelligence and policy to enable security and other services	5

1. Introduction

Encrypting DNS traffic has been a focus of the IETF for several years, and in late 2018 two standards were formalized for use between clients (stub resolvers) and resolvers¹: DNS over TLS and DNS over HTTPS. Numerous implementations have appeared, and DNS encryption has become a visible topic in industry media.

It's a testament to the original design that the way the DNS operates has remained largely unchanged for more than 30 years since the protocol was originally specified. Stub resolvers on clients (typically configured from a local network with a protocol like DHCP) send queries to a caching resolver which, in turn, talks to authoritative DNS servers that provide answers to queries.

DNS encryption changes the transport protocols and, due to some design choices, opens up the possibility of significant changes in the way client devices behave. This paper discusses these changes and their potential impact on service providers. It also offers guidance about how to address encrypted DNS deployments, summarized below:

- Communicate about privacy and security practices so subscribers are aware of how their service is protected and privacy is preserved
- Implement Best Practices for DNS resolution to ensure services are performant, resilient, and always available
- Understand the new DNS encryption protocols and how they can be deployed, and participate in formulation of standards to ensure they can be scaled and operationalized
- Consider additional services that protect subscribers and further enhance their privacy by preventing loss of personal data

2. DNS Encryption Protocols

The DNS over TLS protocol (DoT) is specified in IETF RFC 7858. DNS over TLS uses port 853 rather than port 53 originally specified for DNS. Currently available client implementations of DoT are summarized in a table below. It's important to point out stub resolvers on user devices can also connect to “over the top” public DNS services rather than an in-network resolver provisioned by a network operator². Because it uses a dedicated port, it is easy to detect DoT in network traffic, a useful characteristic for network operators and security teams.

¹ Details of the motivations for developing these new protocols can be found on the Akamai Blog: [Architectural paths for evolving the DNS](#)

² Public DNS resolvers have been available for many years but the advent of DNS encryption creates the perception they are more private and secure, and integration into client software makes them more accessible.




Figure 1 – DNS over TLS defines encrypted transport between stub resolvers and resolvers using TLS

The DNS over HTTPS protocol (DoH) is specified in IETF RFC 8484. DoH uses the same port, 443, as HTTPS. Currently available client implementations of DoH are summarized in a table below. As with DoT, user devices can connect to “over the top” public DNS services. Because it uses the same port as HTTPS, it’s impossible to identify DoH in standard web traffic, which raises obvious security and operational concerns. Perhaps also obvious but worth stating, operators of DoH resolvers still see queries in the clear, regardless of the encrypted transport and services provided by the resolver function as they would with unencrypted transport.

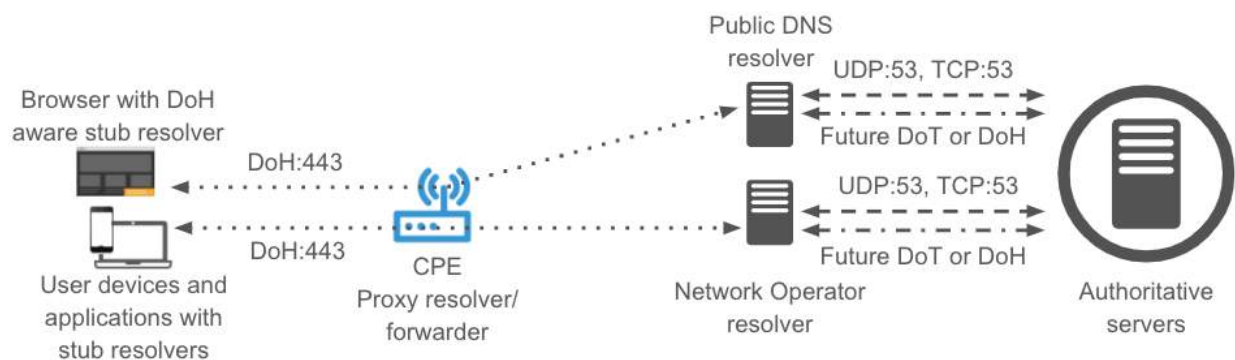


Figure 2 – DNS over HTTPS defines encrypted transport between stub resolvers and resolvers using HTTPS

3. DNS Encryption and ISPs

DNS encryption is intended to protect users from unwanted eavesdropping of DNS traffic by a third party on the path between the user and the resolver they’re connected to. Most provider networks are highly secure, and it’s challenging for adversaries to infiltrate them and intercept traffic. Providers in many parts of the world are also subject to data privacy regulations and/or have contractually agreed Terms of Service that spell out how they use and protect customer data.

This clouds the DNS encryption value proposition for providers. It’s hard to make a business case that secure networks with defensible data protection processes and policies benefit from a layer of encryption that adds cost and complexity.

But service providers are still motivated to understand these new protocols because they may fundamentally change the way subscribers perceive DNS, and client implementations may make it easier for users to bypass provider DNS and connect to public DNS resolution services like Google and several others.

On the positive side, providers have the potential to create value added services that take advantage of DNS encryption. Technology solutions can be developed that keep business and consumer subscribers connected to encrypted DNS resolvers offered by their provider when they're off that provider's network, visiting an untrusted Wi-Fi hotspot for instance. In these cases, since their traffic is transiting untrusted networks, encryption is useful.




Figure 3 – Encrypted transport termination with query processing in the clear

Queries are de-encrypted at the transport layer and presented in the clear for resolution. As encryption between client and resolver is terminated at the providers DoH/DoT service, the DNS service can be equipped with threat intelligence and policies that identify malicious or unwanted domain names to enable security (blocking phishing and malware for subscribers) or content filtering (parental controls for families) to add more value. Integrated offers can also be created with a unified subscriber experience across both a providers network and other networks a subscriber traverses.

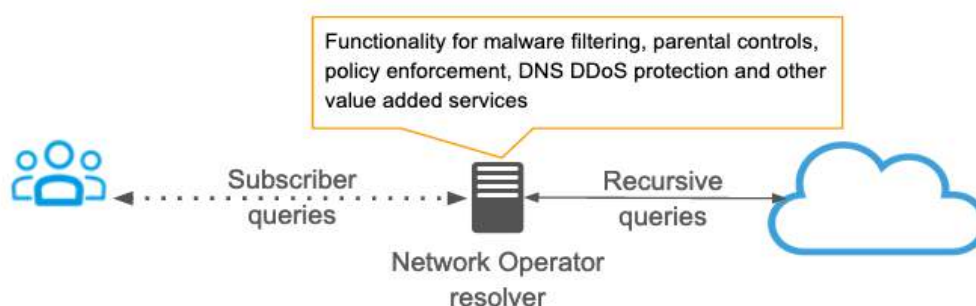


Figure 4 – Resolvers can be equipped with threat intelligence and policy to enable security and other services

4. DNS Encryption Client Implementations

Numerous client implementations of DoT and DoH are tabulated below. The client ecosystem continues to progress at rapid pace, with a broad representation of DNS Encryption capable operating systems, browsers, applications and CPEs existing today. It is expected that configuration mechanisms and awareness of local network conditions, as detailed below, will continue to evolve with ongoing standardization efforts.

Client Feature Summary as of Q3 2020

	DoT	DoH	Existing Configuration Mechanisms	Awareness of local network conditions
Operating Systems				
Android 9+			Same Provider Auto Upgrade or user specified on OS config	Auto upgrade to same DNS provider, fallback to unencrypted
Apple iOS 14			Configurable by Apps or user specified on OS/system config	Enterprise policy awareness. Fail open when auto-discovery in use. Resolver configuration specified by App is optional fallback to system specified encrypted DNS resolver
Apple MacOS 11			Configurable by Apps or user specified on OS/system config	As above
Windows 10			Limited Same Provider Auto Upgrade or user specified on OS/system config	User specified (Unencrypted only, encrypted only, encrypted preferred with unencrypted allowed)
Browsers				
Firefox			Geo Specific Opt-Out + explicitly configured	Canary domain. Fallback to system specified DNS. Enterprise policy, safe search, parental controls detection
Chrome			Limited Same Provider Auto Upgrade + explicitly configured	Auto upgrade to same DNS provider, fallback to unencrypted. Enterprise policy, parental controls detection
Chromium variants			Limited Same Provider Auto Upgrade + explicitly configured	As above
Mobile Apps				
1.1.1.1			Manually enabled, restricted to 1.1.1.1 service	User specified App exclusion
Intra			Manually enabled	User specified App exclusion

Quad9 Connect			Manually enabled, restricted to 9.9.9.9 service	User specific list of domains to send to system resolver
CPE				
FritzBox			Manually enabled	User specified (DoT servers, fallback behavior)
Turris			Manually enabled	User specified (DoT servers)
OpenWRT			Manually enabled	User specified (DoT/DoH servers)

Additional details of current client implementations can be found in the appendix at the end of this paper.

5. Provider Impact of DNS Encryption Clients

As can be seen from the descriptions above, there's currently considerable diversity in client behavior because standards only define how to use secure transport for DNS. Standards aren't yet defined for clients to discover encrypted resolvers, understand local network conditions, and establish and maintain a connection.

Today end users need to take some action in order to enable DNS encryption - navigate to a configuration interface and accept defaults and/or enter information or load an app. There are also differences in how clients fall back to DNS over port 53 if a connection to an encrypted resolver can't be established or fails. And there's no agreed upon method to acknowledge or detect local network conditions, such as the presence of a VPN, an enterprise network, or DNS filtering that might be subverted by the choice of an alternative resolver.

These are critical limitations for service providers. Manual configuration by users is completely incompatible with operation at scale. Default configurations that favor public DNS resolvers bypass provider DNS. Ignoring local network conditions can subvert security and services like parental controls.

As of July 2020, a wide range of possible solutions to these problems have been proposed in the IETF. They can be broadly categorized as: informational drafts describing the current state of the problem, proposals to use existing network technologies like DHCP or Radius to upgrade to secure transport, methods to add functions to the DNS itself, and overlay solutions.

One of the drafts is currently being tested with the Firefox browser and DoH resolvers deployed by Comcast³. In simplified terms, it tests for the presence of DNS policy (e.g., security, parental controls) using a "canary" domain that signals its presence and then queries for a special name to get the address of an encrypted resolver provisioned on the local network. If the query fails, then additional logic can be implemented to select an alternative DoH resolver. This mechanism is currently being tested in Firefox with DoH resolvers deployed by Comcast.

³ <https://www.ietf.org/id/draft-rescorla-doh-cdisco-00.html>

To influence the way clients discover resolvers, the ISP/MNO community needs to be active in the IETF and contribute to relevant RFCs. This will ensure standards deliver the same “just works” experience users have today and are compatible with operational systems.

6. Operational Impact of DNS Encryption

Privacy is a highly visible issue almost everywhere in the world and, if subscribers perceive encrypted DNS is “better,” providers may be motivated to deploy it across their resolution infrastructure, particularly if subscribers start to migrate toward public DNS services.

Providers need to consider underlying details of client implementations because they’ll impact operation and scaling of DNS resolution infrastructure. The shift to TCP-based, secure, transport is a major change from almost exclusively UDP-based transport today. To maximize network efficiency, resolvers will have to support today’s TCP and UDP as well as multiple transport-level authentication and encryption options going forward. Dedicated equipment for TLS termination (like load balancers) increases costs and operational burden. It also adds complexity to troubleshooting efforts with separate interfaces for transport layer problems and DNS resolution itself. Different operational teams need to be coordinated to resolve issues.

Resolver performance will be heavily driven by client side implementations, and there’s little consistency at present. Connection set up overhead must be understood and resolvers need to be tuned for TCP based services in addition to UDP based services. Session reuse (reusing established sessions for multiple queries) must be evaluated as well since it can have a large impact on performance. Advancements in modern server hardware allow for comparable scaling of TLS termination negating the need for specialized appliance based solutions. Failure conditions and resultant bursts of connection setup requests also need to be factored into dimensioning decisions.

Traffic types will shift as client implementations change. Monitoring and comparing Do53, DoT and DoH workloads on an ongoing basis will allow operations teams to make educated capacity planning decisions. Insights into these factors can be found in a presentation at the DNS-OARC conference held in February 2020: [DNS Encryption Operational Experience and Insights](#).

Because it runs over HTTPS, the advent of DoH introduced the possibility of tighter integration with applications, and DoH implementations have been released in browsers. Any app could choose to implement DoH and it is also supported in several public or “over the top” DNS resolution services. The combination of these two developments has important implications for ISPs (and other network operators) and the people who use their networks.

Migrating DNS resolution to applications is a significant change. In the past, applications running on devices relied on a stub resolver implemented as part of the devices operating system which typically query resolvers provisioned by the operator of the network a device is connected to.⁴

⁴ Most operating systems also allow users to manually configure DNS settings to point to a resolver that will take precedence over a resolver configured by the local network.

Fragmentation of DNS resolution among applications raises a number of concerns. One of the most obvious is the risk of substantially complicating troubleshooting when connectivity problems arise. As can be seen from the table and appendix there is currently considerable diversity in client implementations. Individual applications could choose different resolvers and have different methods for exposing that choice to the user (or not expose it at all). They could also have different philosophies about respecting local DNS filtering policies on a network.

Additional considerations may apply when a provider offers resolution services to their enterprise customers. Businesses are potentially exposed when workers use public DNS services, knowingly or not, because sensitive internal domain names could be leaked to external sources. Internal enterprise applications also will not work properly since internal names will not resolve on public resolvers. Enterprises or provider partners may need to make provisions to block access to public DNS services to prevent these problems.

Providers may need to adjust services that use DNS filtering for parental controls or security protections to account for the presence of client implementations that may choose public DNS services. For example as discussed previously some client implementations attempt to check for the presence of DNS filtering by querying for a special canary domain name. In order to ensure their filtering services are preferred by these clients providers will need to provision canary names and respond to the queries properly. In the future there may be other methods that will have to be accommodated.

It appears as though the threat landscape will evolve as well, as attackers explore whether an encrypted DNS path offers advantages.⁵ In the past it was easy to monitor DNS traffic but encrypting the transport with DoH complicates the picture since DNS queries look no different than massive volumes of HTTPS traffic traversing a network.

One possible solution is to break the bootstrapping mechanism exploits use. DoH stub resolvers have to query special hostnames to obtain the IP addresses of DNS resolution services before they can establish a connection. The stub has to use the default resolver in the operating system on the device where it resides in order to accomplish this, which is usually configured by the local network (such as a provider network). Security vendors can track host names of malicious third party DoH resolvers so access to them can be blocked.

Regardless of whether DNS encryption is deployed, the presence of public alternatives amplifies provider incentives to ensure their DNS resolution services are robust and performant. Resolvers are the glue that connects subscribers to their fixed and mobile broadband services. If operators of public DNS services succeed in persuading subscribers to use their resolvers, they will play a significant role in controlling the user experience. DNS is central to virtually every internet transaction, even simple web page loads can send tens of queries. This means performance and latency of public DNS resolvers can have a direct impact on how a user perceives their internet access. Although relatively rare, there may also be cases where a resolver is unable to resolve a name and users get an error message that a resource is unavailable.

When public DNS services operate slowly or fail, as has happened several times in the past, subscribers may associate the problem with their service provider because they may not understand the role DNS plays or may not remember they switched their DNS settings! Providers may need to bear support costs during 3rd party outages and deal with unhappy customers.

⁵ <https://www.sans.org/reading-room/whitepapers/dns/needle-haystack-detecting-dns-https-usage-39160>

A blog post referenced in the bibliography offers design and deployment guidance to help providers establish their resolvers as the preferred choice.

7. Summary Action Plan

DNS encryption has been highly visible in industry media for more than 2 years; there are many client implementations available including those from major OS, browser, mobile app and CPE vendors and several scaled public DNS resolution services exist that support it. Whether deploying it or not, providers need to be aware of the landscape and prepared to respond by:

- Communicating about privacy and security practices including network protections that block intruders, DNS data usage and retention policies, and other privacy enhancing measures in place.
- Implementing Best Practices for DNS resolution whether or not DNS encryption is supported to ensure provider resolvers are better than OTT alternatives - more responsive, reliable, resilient, & secure.
- Considering value added services that protect subscribers by deterring phishing, bots and malware that invade privacy and steal valuable personal data. Motivate subscribers to personalize their service - so they're less likely to leave.
- Contributing to relevant standards to ensure DNS encryption implementations deliver the same "just works" experience users have today, and are compatible with operational systems.

Abbreviations

CPE	Customer Premises Equipment
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
DoT	DNS over TLS
DoH	DNS over HTTPS
HTTPS	Hypertext Transfer Protocol Secure
IETF	Internet Engineering Task Force
ISBE	International Society of Broadband Experts
OTT	Over The Top
OS	Operating System
RFC	Request For Comment
SCTE	Society of Cable Telecommunications Engineers
TCP	Transport Control Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol
VPN	Virtual Private Network

Appendix: DoT and DoH Client Implementations

As of August 2020, implementations will continue to expand and evolve.

DoT

Operating systems

Google Android - first with support for a DoT client in 2018. After configuration by the user (it's not a default yet) it acts as the DNS client for the device, just like the DNS over port 53 client.

Apple iOS 14 and MacOS 11 - Introduced at 2020 developer conference. There is currently no way to configure DoH/DoT from the network. Users can configure an encrypted default resolver for all apps on the system. App developers can configure an encrypted resolver independent of the system, and allow users to opt-in or configure their own encrypted resolver. DoH and DoT are context-aware, when a VPN app or corporate network is detected they will not override configured settings. Developers can also write "rules" to enable encrypted DNS in certain situations or contexts. Enterprise administrators will be able to use Mobile Device Management to configure or override encrypted DNS settings. Plans call for warning users if network providers block encrypted DNS.

<https://www.zdnet.com/article/apple-adds-support-for-encrypted-dns-doh-and-dot/>

Mobile Apps

1.1.1.1 - a special purpose app released by Cloudflare in 2019 acts as the default stub resolver for a device. It connects to Cloudflare's 1.1.1.1 public DNS service using DoT or DoH.

Quad 9 Connect - a special purpose app for Android and iOS released by Quad 9 in 2019 acts as the default stub resolver for a device. It connects to Quad 9's public DNS service using DoT. Quad 9 is a nonprofit founded by IBM, Packet Clearinghouse, and the Global Cyber Alliance.

CPE

FritzBox, Turris, and OpenWRT implementations proxy client requests coming in from port 53 over a secure port 853 to a resolver.

DoH

Operating Systems

Apple iOS 14 and MacOS 11 - as above for DoT

Microsoft Windows 10 - a testable version of DoH was released in May 2020 aimed at power users. Users configure Windows to use DoH and then need to separately add encrypted resolvers from Google, Cloudflare, or Quad 9 through the Control Panel or Settings.

<https://techcommunity.microsoft.com/t5/networking-blog/windows-insiders-can-now-test-dns-over-https/ba-p/1381282>

Microsoft plans to release additional features in their 21H1 update to include more accessible system wide configuration functionality

<https://www.howtogeek.com/685996/whats-new-in-windows-10s-21h1-update-coming-spring-2021/>

Browsers

Mozilla - early entrant with experimental Firefox release supporting DoH in June 2018. Currently Firefox falls back from DoH to operating system defaults for DNS when heuristics detect an enterprise DNS configuration or DNS-based parental controls. One of the heuristics is the use of a canary domain, a special domain name implemented by a network operator Firefox can query that signals the use of DNS filtering on a network.

<https://support.mozilla.org/en-US/kb/firefox-dns-over-https>

In 2018 they released requirements for Trusted Recursive Resolvers (TRR) organizations must meet if they want their DoH services accessible in Firefox.

<https://wiki.mozilla.org/Security/DOH-resolver-policy>

Chrome - early entrant with experimental Chrome release in July 2019 . Chrome preserves the user experience by doing Same Provider Auto Upgrade (auto-upgrading when the existing DNS provider supports DoH). It will also allow manual config of a 3rd party DoH resolver.

<https://www.chromium.org/developers/dns-over-https>

<https://blog.chromium.org/2020/09/a-safer-and-more-private-browsing.html>

For completeness Bromite, Brave, Edge, Opera, and Vivaldi all take advantage of DoH features built into Chromium. Their network characteristics are like Chrome.

Mobile Apps

Intra - a special purpose app released by Google's Jigsaw technology incubator in 2019 acts as the default stub resolver for a device. It connects to Google Public DNS using DoH.

<https://getintra.org/#/>

1.1.1.1 - as above for DoT.

Bibliography & References

RFC 7858 Specification for DNS over Transport Layer Security (TLS)

<https://tools.ietf.org/html/rfc7858>

RFC 8484 DNS Queries over HTTPS (DoH)

<https://tools.ietf.org/html/rfc8484>

Akamai Blog - Architectural paths for evolving the DNS

<https://blogs.akamai.com/2018/10/architectural-paths-for-evolving-the-dns.html>

Akamai Blog - Smart DNS: Delivering the Best Subscriber Experience

Use search, link not assigned at time of publication

Dynamic IUC for OFDMA Transmission

A Technical Paper prepared for SCTE•ISBE by

Hongbiao Zhang

Architect, Wireless Solutions
Casa Systems
100 Old River Rd
Andover, MA01810
978 688 6706 x 6462
hongbiao.zhang@casa-systems.com

Peter Wolff

VP, Wireline Solutions
Casa Systems
100 Old River Rd
Andover, MA01810
978 688 6706 x 6403
peter.wolff@casa-systems.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Background.....	4
2.1. Performance Measurement	4
2.2. OFDMA Profile Management.....	5
3. Per-Minislots Measurements.....	5
3.1. OUDP Test.....	6
4. Dynamic IUC.....	7
5. Experimental Results	8
6. Miscellaneous Notes.....	10
6.1. Dynamic Ranging Zone	10
6.2. Partitioning CMs to CM Groups.....	10
6.3. Single-Minislots Approximation.....	11
7. Conclusion	11
Abbreviations.....	11
Bibliography & References	12
Acknowledgements	12

List of Figures

Title	Page Number
Figure 1 A Snapshot of Upstream Interferences in an HFC Network	3
Figure 2 Designated Minislots For Single-Minislots Measurements	6
Figure 3 Dynamic Scheduling IUC	7
Figure 4 Per-Minislots Measurements and DS-IUC Without Noise	9
Figure 5 Per-Minislots Measurements and DS-IUC with Noise	10

1. Introduction

In a cable plant, it is essential to manage upstream spectrum and mitigate the impact of interferences, especially for lower spectrum bands that are susceptible to various types of noise ingress. The introduction of OFDMA in DOCSIS 3.1 (D3.1) with different bit-loadings at different minislots provides the benefits of enhanced capacity, as well as all flexibility in managing the spectrum usage. Yet surprisingly, not many MSOs have taken full advantage of DOCSIS 3.1 capabilities. Part of the reason is, accompanying the enhancements in capacity and flexibility, comes the complexity in computation, as we have to deal with finer granularity in both detecting the noise levels and reacting to them fast enough.

D3.1 provides an OFDMA profile, which defines, among other properties, a bit-loading pattern that could be adopted by a group of cable modems (CMs). Further, a Profile Management Application (PMA) utilizes the power of offline servers to tackle the computation complexity. A PMA server may take performance measurements on an OFDMA channel for a considerably long time in order to calculate OFDMA profiles. However, in a typical cable plant, especially in its lower band, we may observe many random bursts of noise that vary in time, frequency and power amplitude, and such noise might come and go swiftly. Figure 1 below gives a snapshot of upstream interference in an HFC network. Therefore, using PMA alone is not enough. As upstream interference is naturally observed at the CMTS, it makes sense to augment PMA with a CMTS-based solution that adapts quickly to interference levels measured locally.




Figure 1 A Snapshot of Upstream Interferences in an HFC Network

This paper presents a novel approach that uses various OFDMA profiles with different bit-loading configurations to generate a dynamic profile. We consider upstream impairment with a duty cycle of one or a few seconds, and target a CMTS-based adaptive solution that yields optimal throughput across all CMs. To achieve the objective, we measure upstream channel impairments at a per-minislot level through constantly monitoring the receptions at the CMTS burst receiver. Based on the results, we could dynamically upgrade or downgrade IUCs. The decisions to upgrade or downgrade IUCs are considered for each individual minislot, and can take place locally and automatically without notifying the affected modems.

This approach is referred to as “Dynamic IUC” in the rest of the paper. The generated dynamic profile is referred to as “Dynamic scheduling IUC”, or “DS-IUC” in short.

This approach could be combined with any existing profile management mechanisms. For example, an operator could manually configure 2 OFDMA profiles on an OFDMA channel, or alternatively, a PMA server could elect 2 such profiles based on its calculation. The CMTS could then utilize these 2 profiles, plus NULL bit-loading as a special case, to generate a DS-IUC, thus providing a 3-tier adaptation for each minislot, at an interval much shorter than that of the PMA updates.

2. Background

OFDMA transmission as introduced in D3.1 utilizes different bit-loadings at different subcarriers in the same upstream channel, so that it can adapt to different noise conditions at these subcarriers. D3.1 introduces multiple OFDMA Upstream Data Profiles (OUDPs) on an OFDMA channel, each of which defines a bit-loading pattern that could be adopted by a group of CMs at a certain time period. In particular, an OUDP includes the following information:

- An IUC number (5, 6, 9, 10, 11, 12, or 13)
- Bit-loading and pilot pattern for each minislot or consecutive minislots in order. The bit-loading number ranges from 0 (no transmission) to 12 (4096-QAM)

Per D3.1 MULPI ([1]), the content of these profiles is communicated to CMs through UCD messages. The assignment of one or two such profiles to a CM is using TCC in either registration response or DBC messages. The standard limits that a maximum of two profiles of a channel could be assigned to a CM at any given time.

For terminology, as this document considers data transmission on OFDMA channels for the most part, therefore without mentioning explicitly, we may simply use “OFDMA profile”, “data profile” or even “profile” to replace the full term of “OFDMA Upstream Data Profile”. We also use it interchangeably with “data IUC” or “IUC” without ambiguity.

2.1. Performance Measurement

There exist multiple mechanisms to measure the performance of an upstream channel, including:

- Active and Quiet Probe (see [3]). For this purpose the CMTS may schedule one or more OFDMA symbols and sends a P-MAP. In case of an Active Probe, the CMTS uses a SID assigned to an active CM, and measures RxMER of the CM at the time specified by the P-MAP. The active probe symbol for this capture normally includes all non-excluded subcarriers across the OFDMA channel. In case of a Quiet Probe, the CMTS uses an Idle SID in the P-MAP. It requires all subcarriers, including excluded ones, to be probed. Quiet probes could be used as a baseline to generate initial upstream profiles before modems come online. On the other hand, Active Probes could be scheduled periodically for each active CM in order to generate more accurate upstream profiles over time.
- MER and FEC measurement associated with data receptions. The CMTS burst receiver may provide the following information for each received burst:
 - o MER
 - o FEC correctable codeword count
 - o FEC uncorrectable codeword count
 - o Total codeword count

With the codeword counts the CMTS could subsequently calculate FEC correctable rate (i.e., cFEC) and FEC uncorrectable rate (i.e., uFEC).

- OUDP Test (see [1]). D3.1 includes a mechanism to test an OUDP profile currently in use by providing grants with a special OUDP Test SID to the testing CM. The CM will transmit using a specific payload pattern and the CMTS could subsequently count the FEC and CRC errors in addition to measuring the MER. The OUDP Test SID has to be assigned to the CM in advance.

In this paper, we assume the above performance measurement mechanisms are available as a prerequisite. On top of that, we introduce a novel method that derives performance metrics per-minislot and per CM group, in real-time and with no or minimum overhead in bandwidth utilization.

2.2. OFDMA Profile Management

The specifications in reference [4] provide operators a way to statically configure a set of OFDMA profiles on an OFDMA channel. With such a configuration, a CMTS could locally decide which profiles to be assigned to which CMs and at what time.

Reference [2] describes an architecture using an external Profile Management Application (PMA) server. The PMA server constantly monitors the upstream spectrum, by initiating RxMER measurement for upstream subcarriers (using quiet or active probes), or triggering “OUDP Test”, both through the CMTS. The PMA could also utilize information from other tests, such as upstream captures from a PNM server, as well as historical information obtained on the plant. With such information, the PMA server is able to evaluate the current upstream channel’s performance and generate a set of profiles for the channel. Meanwhile, the PMA server may also designate one or two of these profiles to each CM on this channel. Reference [2] defines an API for this operation.

The above mechanisms require assignments or re-assignments of IUCs to CMs, or modification of IUC content over time. Note that there are only a limited number of IUCs that can be supported on an OFDMA channel. Besides, only one or two IUCs on this channel can be assigned to each CM at any given time. One of them is typically the lowest bit-loading IUC (IUC 13) that is required for pre-registration. Therefore, the mechanisms above require the CMTS to communicate with the corresponding CMs constantly using UCD and/or DBC messages. This imposes a limitation as for how fast the profile management functions could run, and how soon the CMTS system could adapt to channel impairment. An operator may define a time interval, which determines how often to re-evaluate the current profiles and potentially modify or re-assign these profiles to CMs.

In this paper, we assume that a profile management mechanism such as stated above is available as a prerequisite, so that at any given time, each CM will be assigned with a high profile and a low profile, denoted as IUC H and IUC L respectively. We loosely define the interval of profile modification or re-assignment as “profile management interval”. On top of that, we introduce a novel method that dynamically upgrades or downgrades IUCs on an OFDMA channel, at per-minislot granularity and in reaction to real-time noise much faster than the “profile-management-interval”.

3. Per-Minislot Measurements

Assume the CMTS burst receiver is able to collect FEC correctable count, FEC uncorrectable count, total word count and MER for each received burst (see 2.1). When FEC errors are detected with the received burst and the burst size is more than a minislot, it is impossible to determine which minislot has contributed to the impairment due to the effect of interleaving. In order to obtain these metrics on per minislot granularity, we designate some minislots in each OFDMA frame to be used for one-minislot grants, with a pre-defined percentage and with their positions rotate in each frame. See Figure 2 for an illustration. In this way each minislot will be scanned over time with the same frequency. For example, if we designate 1% of the minislots in an OFDMA frame to be used for single minislot grants, it takes 100

frames to scan through every minislot. Depending on the frame size (6 – 36 symbols) and Cyclic Prefix (CP) values, there could be ~1000 - ~8000 frames in a second. If for example there are 3000 frames in a second, any particular minislot will be scanned 30 times in a second, including single-minislot grants at this location assigned to any CM group.

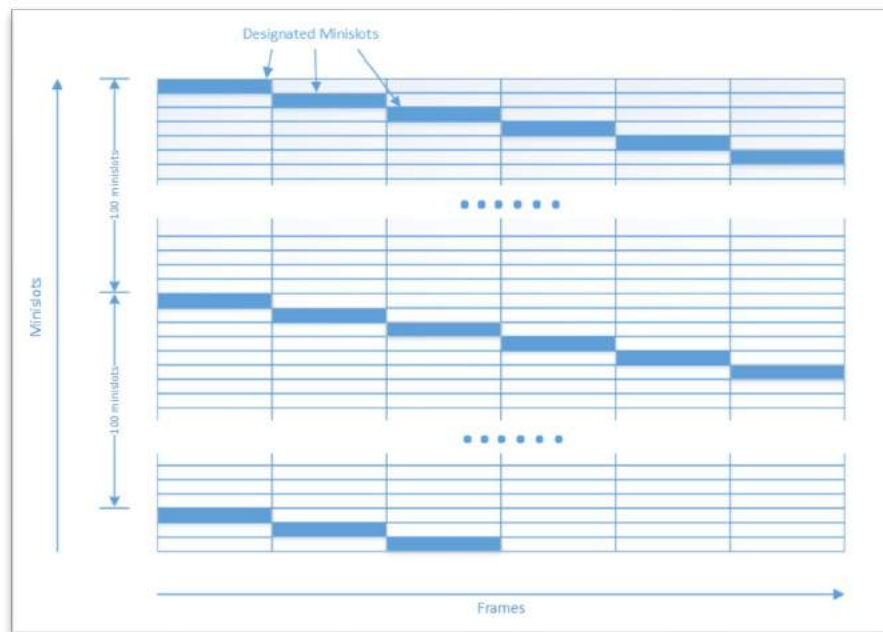


Figure 2 Designated Minislots For Single-Minislot Measurements

At the time of MAP generation, the upstream scheduler allocates grants for CM requests per the normal operation. If a grant hits one of the designated minislot location, this grant will be fragmented, and the next grant will occupy a single minislot. Then, only performance metrics of these single minislot bursts are taken into consideration, and the rest are discarded for this measurement. To be specific, when the single minislot burst is received by the CMTS, it will be measured and the performance metrics will be counted toward the group which the transmitting CM belongs to.

3.1. OUDP Test

There're several scenarios where OUDP Test could be useful, including the following:

- A CM group could possibly drop its bit-loading to NULL in some minislots, due to severe interference experienced at these minislots (see Section 4). In this case the scheduler avoids scheduling any grant for either data or ranging requests, therefore there is no subsequent measurement available on these minislots based on active data. The CMTS then has to explicitly poll these minislots using OUDP Tests, in order to determine whether the interference condition has changed.
- When there are unused minislots in an OFDMA frame, the CMTS could utilize these minislots to test the performance of any CM group of its choice, again using OUDP Tests.

4. Dynamic IUC

Assume with one or more of the “profile management” mechanisms as described in 2.2, a CM is assigned with a high profile and a low profile, denoted as IUC H and IUC L respectively. Then, in a time interval much smaller than the “profile management interval”, the CMTS might switch between IUC H, IUC L, plus NULL bit-loading as a special case. Such dynamic switches could be determined for each individual minislot, and could occur locally without communicating to the CM through MMM messages.

In particular, the CMTS upstream scheduler periodically generates a “Dynamic Scheduling IUC (DS-IUC)”, which adopts IUC H, or IUC L, or NULL, on a minislot-by-minislot basis. In other words, the bit-loading value at a minislot could assume $\text{bit-loading}_{\langle \text{IUC H, minislot} \rangle}$, or $\text{bit-loading}_{\langle \text{IUC L, minislot} \rangle}$, or 0. The scheduler could possibly encode different data IEs of the same CM with different IUC numbers if these IEs fall into different minislots in a MAP.




Figure 3 Dynamic Scheduling IUC

An example of DS-IUC (partial) is illustrated in Figure 3. In this example, IUC H is defined as follows:

{minislot 1: 8-bit QAM}, {minislot 2-4: 7-bit QAM}, {minislot 5-8: 8-bit QAM}, ...

IUC L is defined as:

{minislot 1: 5-bit QAM}, {minislot 2-4: 4-bit QAM}, {minislot 5-8: 5-bit QAM}, ...

The DS-IUC adopts IUC H in minislot 1 – 4, but adopts IUC L in minislot 5 and 6, and again adopts NULL bit-loading in minislot 7 - 8. Therefore the DS-IUC is defined as:

{minislot 1: 8-bit QAM}, {minislot 2-4: 7-bit QAM}, {minislot 5-6: 5-bit QAM}, {minislot 7-8: 0-bit QAM}, ...

The collection of IUC H, IUC L and NULL provides a 3-tier option for dynamic adaptation to noise conditions to a certain extent without involving the CM, and such adaptation may be different on different sub-bands. The creation and modification of DS-IUC is performed periodically. This period is referred to as “dynamic scheduling interval”, which should be much smaller than the “profile management interval”. For example, the dynamic scheduling interval can be chosen from sub-second to multiple seconds while the profile management interval can be chosen from sub-minute to multiple minutes or even hours.

5. Experimental Results

To demonstrate the effects of Dynamic IUC, we ran an experiment in the lab with a D3.1 CMTS from Casa Systems and a D3.1 CM from Technicolor. An OFDMA channel is configured with the frequency span of 5 – 85 MHz and with $K = 18$ symbols per OFDMA frame. 5 distinct Data IUCs are configured and allocated to the OFDMA channel, namely IUC 9 – IUC 13.

When the CM is online, two data IUCs, namely IUC 9 and IUC 13, are assigned to the CM based on assessment of the current condition. Here IUC 9 is the “IUC H” and IUC 13 is the “IUC L”: IUC 9 has a constant bit-loading of 1024-QAM across all minislots, and IUC 13 has a constant bit-loading of 64-QAM across all minislots. Between the two IUCs assigned to the CM, the scheduler starts with IUC H for unicast data transmission until there’s an issue.

Profile management interval is configured as 1000 seconds, which means for every 1000 seconds, the system could possibly update the assignment of IUC H and/or IUC L to the CM. This operation is skipped in our experiment. On the other hand, the dynamic scheduling interval is configured as 5 seconds, which means for every 5 seconds, the scheduler possibly switches among IUC H, IUC L and NULL for each minislot and each CM group, based on per-minislot performance measurements of the current 5-second interval.

Figure 4 depicts a snapshot of the OFDMA channel (partial) when there is no noise. In the figure, a) and b) reflect cMER and uMER respectively, c) is the plot of MER, and d) displays the calculated DS-IUC. At this point DS-IUC is the same as IUC 9 across all minislots. As we can see, using IUC 9 the per-minislot cFEC and uFEC are both 0, and the per-minislot MER stays between 43.1- 47.7dB. These results show IUC 9 is working appropriately across the whole channel, therefore DS-IUC is unchanged, i.e., it stays the same as IUC 9.

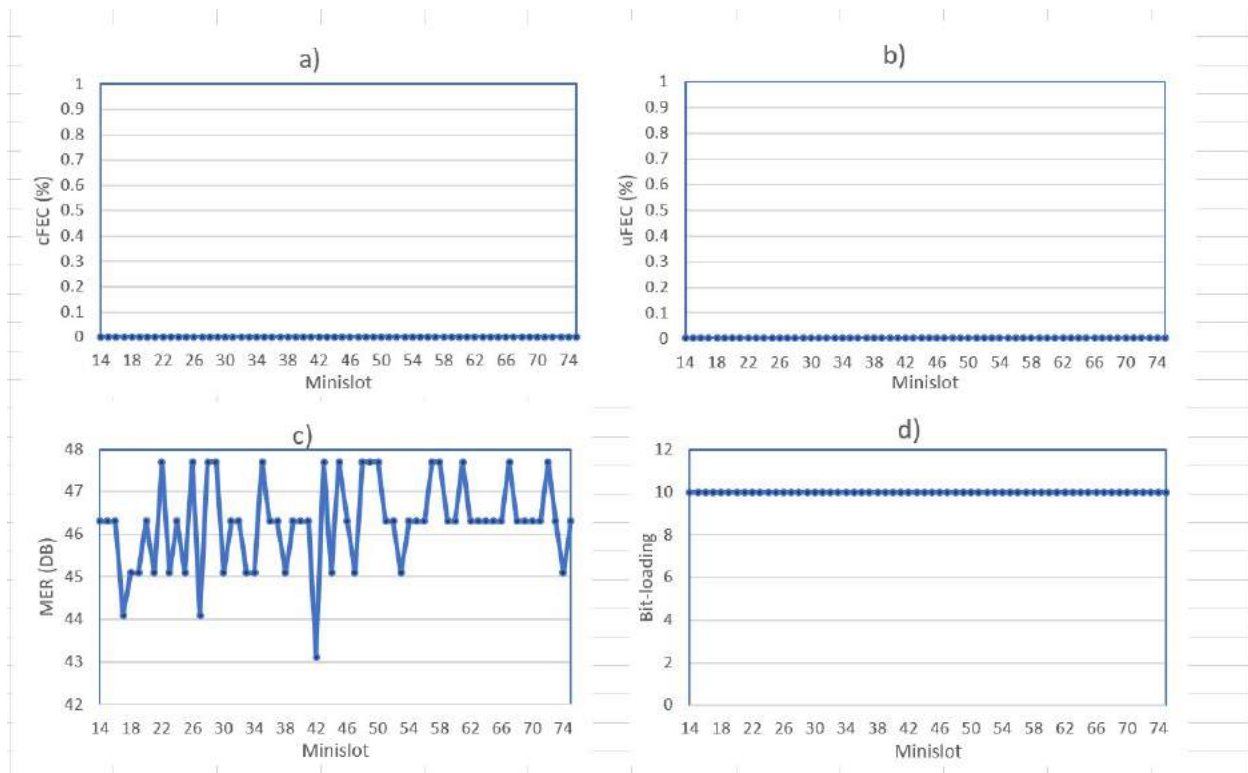


Figure 4 Per-Minislot Measurements and DS-IUC Without Noise

We then introduce some noise using a signal generator. The generated signal is a sharp tone centered at 25Mhz with a FM of 1Mhz. It is injected into the upstream port with a 2-way combiner. Figure 5 depicts a snapshot of the OFDMA channel (partial) roughly 5 seconds after the noise is injected. Again a) and b) reflect cMER and uMER respectively, c) is the plot of MER, and d) displays the calculated DS-IUC. As we can see, at this point the MER drops below 42.3dB within minislot [30 – 68], and drops below 26.7dB within minislot [46 – 53]. This forces a degradation in bit-loadings. As a result, the scheduler adopts IUC 13 for impaired regions within minislot [30 – 45] and within minislot [54 – 68]. It further adopts NULL bit-loading for a severely impaired region within minislot [46 – 53]. The rest of regions is little affected and so IUC 9 is kept unchanged.

Note the diagrams in Figure 5 illustrate measurements after the adapted DS-IUC is calculated and take into effect, which comprises IUC 9, IUC 13 and NULL at unimpaired, impaired, and severely impaired regions respectively. In the impaired regions, since modulation is reduced to IUC 13, uFEC and cFEC become 0 once again. In the severely impaired region however, as NULL bit-loading is adopted, there's no data permitted for transmission. The CMTS needs to schedule one-minislot grants for OUDP Tests, using IUC 13. Diagrams in a) and b) of Figure 5 show that the cFEC and uFEC values obtained with such OUDP Test frames are still significantly high at the severely impaired region, and so the DS-IUC stays at NULL in this region.

When the noise is removed, the per-minislot FEC and MER metrics would change accordingly to what were before. Then, with less than 5 seconds delay, the scheduler adjusts DS-IUC to adopt IUC 9 once again across all minislots. The diagrams at this moment are almost identical to those in Figure 4 and are omitted here.

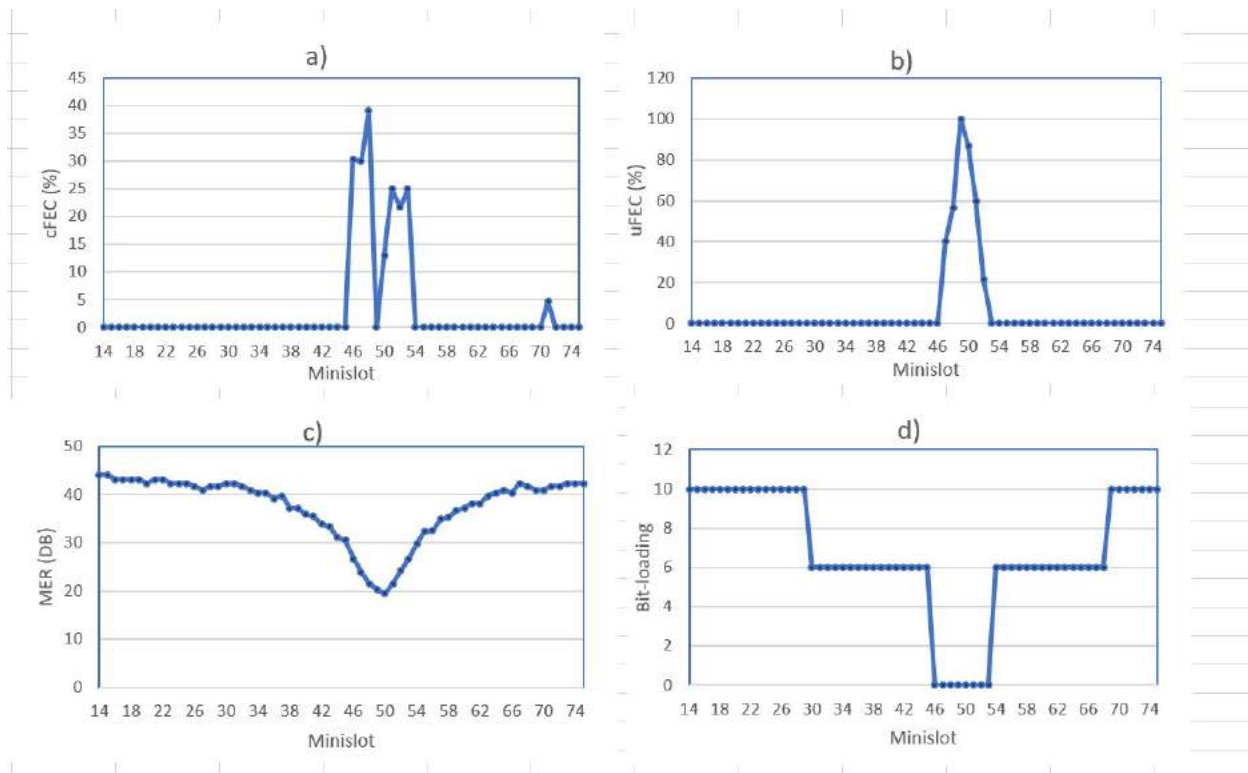


Figure 5 Per-Minislot Measurements and DS-IUC with Noise

6. Miscellaneous Notes

A few notes regarding Dynamic IUC are worth mentioning here, as detailed in the following.

6.1. Dynamic Ranging Zone

Besides utilizing the per minislot performance metrics to facilitate data transmissions, the CMTS could also use the same to direct range requests. For example, the CMTS could choose a default ranging zone to start with. Once cable modems come online and the per-minislot performance metrics are collected, if the CMTS determines that the original ranging zone is too noisy, it could optionally select a new region with the highest MER. It then moves the initial ranging and/or fine ranging zone to the selected region, i.e., the ranging IEs would contain IUC 3 or IUC 4, and contain minislots from within that region.

6.2. Partitioning CMs to CM Groups

In the previous sections we loosely refer to the term “CM group”. It is further explained in this subsection, as follows. We assume that all CMs utilizing an upstream channel could be partitioned into a limited number of groups, with each group observing the same pattern of impairment at almost all times. Different group behaviors could reflect different D3.1 CM types, or different segments of the cable plant, etc.

The methods and rules to partition CMs into CM groups depend on deployment scenarios and implementation specifics. In an extreme case, if the number of CMs is manageable and the resources (memory, cpu cycle) are sufficient, we could designate each CM as a CM group.

In another extreme case, if all CMs on the upstream channel behave uniformly across all minislots and at all times, they could be categorized as in the same CM group. Due to the “funnel effect” of upstream noises, this model could work for lots of scenarios. For example, it could work when the only source of interference is ingress noise.

With CMs on an OFDMA channel partitioned into CM groups, we could obtain FEC and MER statistics on a per CM group basis. We could also define DS-IUCs on a per CM group basis, as described in the previous sections.

The exact mechanism for how to partition CMs into CM groups could be a topic for further studies.

6.3. Single-Minislot Approximation

Special care must be given when a single minislot is not sufficient for a grant. This could happen if the combination of bit-loading and frame size is too small for the smallest LDPC codeword. In that case a smallest multi-minislot grant to fit the LDPC codeword could be scheduled, which covers the minislot being examined plus one or more neighboring minislots. The performance metrics of this minislot will be estimated based on what is measured with the multi-minislot burst. The exact mechanism as for how to obtain the estimation of the single minislot performance metrics could be a topic for further studies.

7. Conclusion

The solution of “Dynamic IUC” presented in this paper starts with selectively scheduling one-minislot grants for either active data or OUDP Test data, and measures channel impairments at a per-minislot granularity using data received at the burst receiver. Based on the above, the CMTS leverages the two OFDMA data profiles, namely IUC H and IUC L, that were previously assigned to a group of CMs, plus NULL bit-loading as a special case, in order to produce a DS-IUC with a 3-tier adaptation. The decision to choose the tier is made for each individual minislot and each CM group. Additionally, the decision takes place locally without notifying the affected CMs. Finally, it can take place immediately, i.e., one or a few seconds after onset/offset of an impairment.

Dynamic IUC could be combined with any existing profile management mechanisms, such as PMA. For example, one could use PMA to elect IUC H and IUC L for a group of CMs, and use Dynamic IUC to leverage the elected IUCs and provide a 3-tier adaptation. In this way the system is able to handle both long-term and short-term impairments, resulting in an optimal upstream throughput across all CMs.

Abbreviations

API	application programming interface
cFEC	correctable FEC
CM	cable modem
CMTS	cable modem termination system
CRC	cyclic redundancy check
D3.1	DOCSIS 3.1
DBC	dynamic bonding change
DOCSIS	data over cable service interface specification
DS-IUC	dynamic scheduling IUC
FEC	forward error correction

IE	Information element
IUC	interval usage code
LDPC	low density parity check
MAP	upstream bandwidth allocation map
MER	modulation error ratio
MMM	MAC management message
OFDMA	orthogonal frequency division multiple access
OU DP	OFDMA upstream data profile
PMA	profile management application
P-MAP	probe MAP
QAM	quadrature amplitude modulation
RxMER	receive MER
SID	service identifier
TCC	transmit channel configuration
UCD	upstream channel descriptor
uFEC	uncorrectable FEC

Bibliography & References

1. *Data-Over-Cable Service Interface Specifications DOCSIS® 3.1, MAC and Upper layer Protocols Interface Specification*, CM-SP-MULPIv3.1-I10-170111
2. *Data-Over-Cable Service Interface Specifications Technical Reports, DOCSIS® 3.1 Profile Management Application Technical Report*, CM-TR-PMA-V01-180530
3. *Data-Over-Cable Service Interface Specifications DOCSIS® 3.1, Physical Layer Specification*, CM-SP-PHYv3.1-I10-170111
4. *Data-Over-Cable Service Interface Specifications DOCSIS® 3.1, CCAP™ Operations Support System Interface Specification*, CM-SP-CCAP-OSSIV3.1-I11-171220
5. White and Sundaresan, *DOCSIS 3.1 Profile Management Application and Algorithms*. SCTE IBSE, NCTA, CableLabs, Spring Technical Forum Proceedings, 2016

Acknowledgements

The authors would like to acknowledge Chain Lee, Yanbo Yuan, Tao Yu, Weidong Chen and Chaoyi Wang from Casa Systems for their valuable contributions. The authors would also thank Don Jones and Charles Moyer from Casa Systems who have provided the interference capture and experimental data respectively.

Full Band Capture Revisited

A Technical Paper prepared for SCTE•ISBE by

Ron Hranac

Technical Marketing Engineer

Cisco Systems

9155 E. Nichols Ave., Ste. 400, Centennial, CO 80112

+1-720-875-1338

rhranacj@cisco.com

and

Chad Campbell, Intraway

Roger Fish, Broadcom

Tom Kolze, Broadcom

Even Kristoffersen, Telia Norge

James Medlock, Akleza

Jason Rupe, CableLabs

Paul Schauer, Comcast

Aleksander Soeberg, Telia Norge

Tom Williams, CableLabs

Larry Wolcott, Comcast

Table of Contents

Title	Page Number
1. Introduction.....	5
2. FBC Benefits.....	7
3. A Closer Look at Impairment Categories	9
4. How FBC Works	10
4.1. General Operation	10
4.2. FBC Controls.....	11
5. FBC Examples.....	12
5.1. FBC Use Cases.....	17
5.1.1. Water damaged drop	18
5.1.2. Localizing a problem to a specific drop	20
5.1.3. FBC Analysis and Fault Location Example	22
6. How to retrieve and Display FBC Data	27
6.1. FBC Data Collection	27
7. Advanced FBC Applications.....	35
7.1. Estimating Distances with High Accuracy.....	35
7.2. Drop Cable Testing.....	37
7.3. Using FBC to View Upstream Noise.....	37
7.4. Using FBC to Find Water-Soaked Coaxial Cable	38
8. Impairment Detection.....	39
8.1. Impairment detection and automation	40
9. Conclusion.....	41
10. Abbreviations.....	41
11. Bibliography & References.....	42
12. Appendix.....	44
12.1. File: getFbcData.py	44
12.2. File: showFbcData.py	45

List of Figures

Title	Page Number
Figure 1. Spectrum analyzer screen shot showing analog TV signals on the left side and digital signals on the right side. The vertical axis is amplitude, and the horizontal axis is frequency.	5
Figure 2. FBC screen shot of a cable network's downstream spectrum from a DOCSIS modem in the home of one of the authors (courtesy of Comcast).....	6
Figure 3. FBC screen shot of a cable network's downstream spectrum. This display shows a suckout (notch) at about 625 MHz (courtesy of Broadcom).....	6
Figure 4. FBC from modem in lab setup (see text). Courtesy of CableLabs.	7
Figure 5. Spectrum analyzer display of the same RF spectrum in Figure 4 (courtesy of CableLabs).	8
Figure 6. Examples of some impairments found using FBC (courtesy of Comcast).....	9
Figure 7. Digital spectrum analyzer block diagram (see text).....	11
Figure 8. Adjacency – While the spectrum is relatively flat, note that signals below 400 MHz are a few dB higher in amplitude than signals above 400 MHz. This suggests incorrect RF level adjustment in the headend or hub site (courtesy of Comcast).	12
Figure 9. Filter – Indicates the presence of a filter in the subscriber drop (courtesy of Comcast).....	13

Figure 10. Negative tilt – Typical response at or near ends-of-line locations, but excessive negative tilt could indicate a problem. Note presence of what appears to be FM broadcast band ingress at the left end of the display (courtesy of Akleza).....	13
Figure 11. Positive tilt – Typical response at or near the output of nodes and amplifiers. Note presence of what appears to be FM broadcast band ingress at the left end of the display (courtesy of Akleza).....	14
Figure 12. Standing wave – Classic example of scalloped sinusoidal wave shape in the response, caused by an impedance mismatch (courtesy of Akleza).....	14
Figure 13. Suckout – Notch in the response, centered just above 700 MHz (courtesy of Akleza).....	15
Figure 14. Suckout - A severe example, centered between 500 MHz and 600 MHz (courtesy of Comcast).....	15
Figure 15. Water damage - Non-periodic wave shape in the response and higher attenuation at higher frequencies, typical of water damage in a subscriber drop (courtesy of Akleza).....	16
Figure 16. Resonant peaking – Day 1 FBC spectrum (courtesy of Comcast).....	16
Figure 17. Resonant peaking - Day 2 FBC spectrum, same modem as previous figure (courtesy of Comcast).....	17
Figure 18. Resonant peaking – Day 3 FBC spectrum, same modem as previous two figures. Note that the problem seems to have disappeared (courtesy of Comcast).....	17
Figure 19. Resonant peaking - Day 4 FBC spectrum, same modem as previous three figures. The response problem has returned (courtesy of Comcast).....	17
Figure 20. Water damage before repair. Note the non-periodic wave shape in the FBC response and the higher attenuation at higher frequencies. This particular example occurred when abrasion damaged the cable's jacket, allowing water to enter the cable (see text). Courtesy of Comcast.....	18
Figure 21. Most of the downstream SC-QAM signals have low signal level and degraded RxMER, indicated in red shaded boxes. The upstream was relatively unaffected (courtesy of Comcast).....	18
Figure 22. Damaged coax jacket where water was able to enter the cable.....	19
Figure 23. Water coming out of the end of the connector at the ground block.....	19
Figure 24. FBC response after drop cable was replaced from the tap to the ground block (courtesy of Comcast).....	20
Figure 25. Signal performance after the drop was replaced (courtesy of Comcast).....	20
Figure 26. Subscriber A's FBC showing 4G/LTE and 5G ingress (courtesy of Telia Norge).....	21
Figure 27. Subscriber B's FBC, showing no 4G/LTE or 5G ingress (courtesy of Telia Norge).....	21
Figure 28. Location of cell site near the affected subscriber drop.....	22
Figure 29. Example of using FBC to isolate a standing wave fault (courtesy of Akleza).....	23
Figure 30. Amplifier [1] containing water and showing corrosion (courtesy of Akleza).....	24
Figure 31. Fault still visible after replacement of amplifier [1] (courtesy of Akleza).....	25
Figure 32. Amplifier [2] with corroded connector (courtesy of Akleza).....	26
Figure 33. FBC trace after faulty connector replaced (courtesy of Akleza).....	26
Figure 34. IF3-MIB spectrum capture objects.....	28
Figure 35. Configuration example for high resolution downstream data collection.....	29
Figure 36. Returned SNMP table row for each frequency span.....	29
Figure 37. FBC plot using Microsoft Excel (courtesy of Akleza).....	31
Figure 38. FBC plot generated using example Python scripts (courtesy of Akleza).....	31
Figure 39. FBC plot with no averaging (courtesy of Akleza).....	32
Figure 40. FBC plot with 32 averages (courtesy of Akleza).....	33
Figure 41. Zoomed view of high resolution and averaged FBC (courtesy of Akleza).....	34

Figure 42. Diagrams showing what is observed at two points on a cable line for the one- and two-reflection cases. The impulse responses on the right can be observed by the FBC and the impulse responses can be observed by an active or passive TDR.	36
Figure 43. Damaged hardline coax that was at one end of an echo tunnel; the other end was a loose seizure screw on a chassis terminator installed in an end-of-line tap. Further complicating things: The damaged cable was submersed in water. Courtesy of CableLabs.....	36
Figure 44. Water-soaked coax, FBC plot on left and impulse response on right. The impulse response is highly dispersed.....	39
Figure 45. Coax with standing wave response. FBC plot is on left and impulse response is on right. The impulse response is not dispersed.....	39

1. Introduction

One of the most versatile pieces of test equipment available to the cable industry is the spectrum analyzer. Spectrum analyzers, which are instruments that display signals in the frequency domain (see **Figure 1**) have for decades been used for headend, hub site, and outside plant maintenance and troubleshooting. However, spectrum analyzers have historically been expensive and not typically available to field personnel; usually not particularly portable; and in some cases, complicated to use.

That said, wouldn't it be nice to have a spectrum analyzer in every cable subscriber's home? What if those in-home spectrum analyzers could be remotely accessed so that technicians didn't have to go to each home to use them? Even better, what if that in-home spectrum analyzer capability already existed and was widely deployed?

Let's answer those three questions in reverse order, starting with the last one first: In-home spectrum analyzer capability does exist and has for several years. Those in-home spectrum analyzers can be remotely accessed, using simple network management protocol (SNMP) or similar. And while not in every cable subscriber's home, they are in a sizeable percentage of homes with Data-Over-Cable Service Interface Specifications (DOCSIS[®]) customer premises equipment (CPE).




Figure 1. Spectrum analyzer screen shot showing analog TV signals on the left side and digital signals on the right side. The vertical axis is amplitude, and the horizontal axis is frequency.

That in-home spectrum analyzer capability is known as full band capture (FBC) and is a cable modem spectrum analysis feature available in DOCSIS 3.0 and 3.1 cable modems. Cable modem spectrum analysis functionality was first defined in version I20 of the DOCSIS 3.0 Operations Support System Interface Specification (CM-SP-OSSv3.0-I20-121113) back in November 2012.¹ See [Ref. 1].²

Figure 2 shows an example of an FBC display of the downstream in a cable network. The data was captured remotely from an FBC-capable DOCSIS modem in the home of one of the authors, showing a problem-free downstream spectrum. The vertical axis is amplitude (dBmV), the horizontal axis is frequency in MHz, and the resolution bandwidth (RBW) is 117 kHz.

¹ The DOCSIS 3.0 specifications were first published in 2006, and the first DOCSIS 3.0 cable modems were certified by CableLabs in 2008.

² References are in the bibliography near the end of the document.




Figure 2. FBC screen shot of a cable network's downstream spectrum from a DOCSIS modem in the home of one of the authors (courtesy of Comcast).

Figure 3 shows another example of an FBC display of the downstream spectrum in a cable network. The display was captured remotely from an FBC-capable DOCSIS modem in a cable subscriber's home, and shows a suckout (notch) in the vicinity of 625 MHz that would otherwise have been difficult to identify without having a technician on-site. As before, the vertical axis is amplitude (dBmV), the horizontal axis is frequency in MHz, and the RBW is 117 kHz.




Figure 3. FBC screen shot of a cable network's downstream spectrum. This display shows a suckout (notch) at about 625 MHz (courtesy of Broadcom).

Some operators have been using FBC successfully as part of their proactive network maintenance (PNM) programs. But many operators have not taken advantage of FBC, perhaps because they don't know about it, or don't know how to use it. Whatever the reason, FBC is a powerful tool that is supported in most DOCSIS 3.0 and all DOCSIS 3.1 cable modems.

This paper highlights the benefits of FBC; provides an overview of how FBC works; shows several examples of the types of impairments that can be identified using FBC, along with some use cases demonstrating how cable operators use FBC to troubleshoot specific problems; and explains how to retrieve and display spectral data from modems (example Python code to retrieve and plot FBC data can be found in the Appendix). Also included is a discussion about using FBC for displaying upstream spectral data – especially noise – a capability supported in some cable modem implementations. Finally, the paper provides guidance and findings about how to transform the visual spectrum data into actionable decisions, including opportunities for automation, and operational expenditure savings.

2. FBC Benefits

As mentioned previously, FBC is being used by several cable operators as part of PNM programs, both in standalone tools and in third party PNM tools. FBC enables operators to have spectrum analysis capabilities wherever FBC-capable DOCSIS 3.0 and 3.1 cable modems have been deployed.

Is FBC really like having a spectrum analyzer in subscribers' homes? **Figure 4** and **Figure 5** show a lab setup comparing the screen shot from an FBC-equipped cable modem with a screen shot from a Keysight spectrum analyzer. Single carrier quadrature amplitude modulation (SC-QAM) signals are located at the low end of the spectrum starting at 261 MHz. Two orthogonal frequency division multiplexing (OFDM) signals are located between about 408 MHz to 600 MHz and 608 MHz to 800 MHz respectively. A continuous wave (CW) carrier is present at 303 MHz. A filter created a deep notch at about 337 MHz, and a 5 MHz to 750 MHz bandpass filter rolled off the higher OFDM signal about 25 dB from 750 MHz to 800 MHz.




Figure 4. FBC from modem in lab setup (see text). Courtesy of CableLabs.



Figure 5. Spectrum analyzer display of the same RF spectrum in Figure 4 (courtesy of CableLabs).

FBC allows operators to obtain remote spectrum captures at the subscriber premises without having to take an expensive spectrum analyzer into the field; without requiring access into subscribers' homes; and without having to roll a truck to identify problems! Key benefits of FBC include significant operational advantages, not to mention operational cost savings.

With a substantial number of FBC-capable DOCSIS cable modems already deployed, these devices have gathered critical mass and are now in a large percentage of subscribers' homes. As such, operators can use FBC to identify individual drop problems, and correlate RF spectrum signatures from neighboring modems to troubleshoot and locate distribution network problems affecting groups of modems.

The width of the FBC spectrum – the full downstream spectrum (and sometimes also the upstream spectrum, depending on modem implementation) – allows operators to detect very short micro-reflections not visible when just a single channel is analyzed through adaptive equalization or pre-equalization. FBC has enabled verification of channel RF level alignment – for instance, incorrect narrowcast injection signal levels – and the detection of ingress. Correlation of FBC signatures before and after actives enables the detection of nonlinear problems in amplifiers. And much, much more. **Figure 6** shows examples of some of the impairments that can be identified using FBC. Additional examples are included later in this paper.

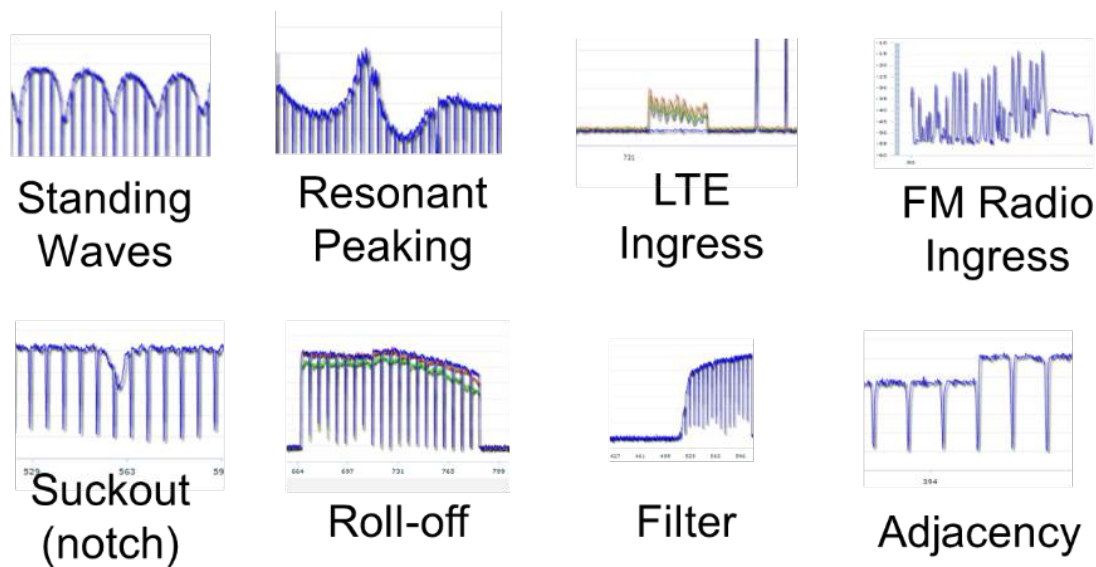


Figure 6. Examples of some impairments found using FBC (courtesy of Comcast).

Clearly, having the equivalent of a spectrum analyzer in every home that has an FBC-capable modem provides powerful troubleshooting and maintenance capabilities. That said, it is simply not practical for cable operators to manually look at every modem's FBC spectrum data. What to do? Automated spectrum signature analysis and impairment detection are available to scale the analysis to the millions of modems deployed with FBC functionality, and are discussed later in Section 8. All of this can be done remotely!

3. A Closer Look at Impairment Categories

Impairments to an FBC's spectrum can be broken into two classes or categories: things that are added to the FBC that should not be there, and downstream signals that have been modified from their original unimpaired state. Fortunately, most downstream SC-QAM and OFDM signals appear to have a relatively static noise-like appearance to a spectrum analyzer, but this is not true for upstream signals.

Energy that is added includes:

- Random noise from plant nonlinear distortion or from too-low RF levels.
- Wideband ingressors such as LTE signals and UHF TV broadcast signals, entering the network via shielding problems in the hardline distribution network or subscriber drop (including inside the home).
- Narrowband signals like broadcast FM radio and amateur radio transmissions, entering the network via shielding problems in the hardline distribution network or subscriber drop (including inside the home).
- House wiring noise. House wiring noise can enter the cable plant through defective coaxial cable shielding, which is unfortunately pervasive. [Ref. 2] estimates that 12% to 20% of houses have defective cable shielding and can pass noise from house electrical wiring to the cable plant. House wiring noise is harder to detect and measure because it is often intermittent, usually impulsive in nature, and concentrated in the upstream band, which is visible only with cable modem FBC hardware that has an upstream view capability. Such modems are already deployed in the field (more on this later),
- FBC measurement artifacts, such as an elevated noise floor and spurious signals. These generally are relatively small and easy to identify.

Things that modify the waveform include:

- Adjacency – One example is incorrect narrowcast RF signal levels compared to the rest of the downstream spectrum.
- Filter – Indicates the presence of a filter (e.g., a high-speed data-only filter).
- Non-periodic waves – A response impairment that occurs in the presence of distributed impedance mismatches that can be caused by, for example, water in the coaxial cable.
- Resonant peaking – Usually caused by a loose module, circuit board, poor grounding, etc., in an active or passive device.
- Rolloff – Low-end rolloff is usually caused by a loose center conductor seizure screw or similar. Higher-frequency rolloff is caused by water, or damaged cable or equipment.
- Standing waves – Periodic scalloped sinusoidal or sinusoidal-like amplitude ripple in the frequency response is created by discrete impedance mismatches.
- Suckout – Notch in the frequency response affecting one or more channels, caused by loose modules, module covers, printed circuit boards, poor grounding, and similar problems inside of active or passive device housings. Can also be caused by repetitive, regularly spaced impedance discontinuities in coaxial cable.
- Tilt – Some upward or downward tilt is normal, but excessive downward tilt is indicative of a problem.

Because a cable operator may have many millions of cable modems in its plant, and because the artifacts can be dynamic, as mentioned previously software exists that can automatically detect and report on both added and modified impairments. The power of human eyes and brains is assisted greatly by sophisticated PNM software algorithms processing thousands of FBC responses in batch mode.

Several examples of FBC spectrum displays showing a variety of impairments are included in Section 5.

4. How FBC Works

4.1. General Operation

FBC is a feature that uses low-cost discrete Fourier transform (DFT) and fast Fourier transform (FFT)-based technology to support spectrum analyzer-like functionality in many DOCSIS 3.0 and all DOCSIS 3.1 cable modems.

As mentioned earlier, FBC spectrum data can be accessed remotely using SNMP or similar, allowing a cable operator to see where various impairments might be problematic from the perspective of the modems and without the need to be on-site at the subscriber premises.

Spectrum analyzers are instruments used to measure the frequency content of an input signal. This is also what DFT does. Multiplying the input signal by what is called a DFT matrix measures the correlation of that input signal with each row in the DFT matrix, and each row is a sine/cosine of a particular frequency. Thus, each output bin represents the power of the input signal at that frequency with a complex value which corresponds to the amplitude and phase of that frequency component within the input signal.

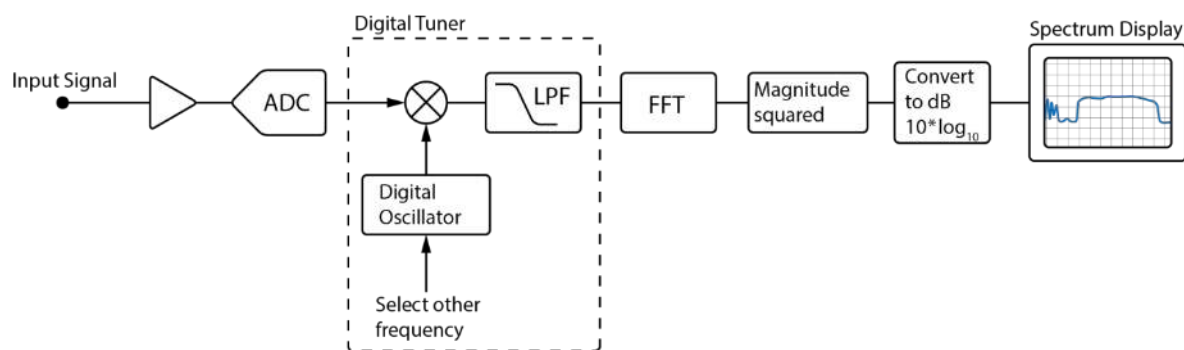


Figure 7. Digital spectrum analyzer block diagram (see text).

Figure 7 shows a block diagram of a digital spectrum analyzer which could reside in a cable modem. The input signal enters at the left of the diagram; this signal is the full upstream or downstream band of the cable plant. An analog front end (represented by a small triangle in the figure) amplifies the signal and provides RF gain control. A high-speed analog-to-digital converter (ADC) provides digital samples of the signal. A digital tuner, consisting of a digital oscillator and lowpass filter, selects the desired analysis band around a specified center frequency. The signal from the selected band is applied to the FFT, which multiplies the signal by the DFT matrix. Each bin of the FFT output comprises a complex value consisting of two numbers, real (I) and imaginary (Q), giving the correlation of the input signal with the particular frequency corresponding to a single row of the DFT matrix. Typically, a spectrum analyzer is only concerned with the magnitude, not the phase, of the FFT output. So, the power (magnitude-squared) of each bin is computed, that is, $I^2 + Q^2$ for each bin. If spectrum smoothing is to be applied, the previously-described process is repeated with a fresh set of data from the same band, and the power values from several captures are averaged at each bin location. The smoothed bins are converted to decibels by taking $10 \cdot \log_{10}$ of each bin power value. The decibel values, one for each frequency bin, are displayed as the spectrum of the input signal.

If the entire band is able to be processed as a single analysis band, the tuner shown in **Figure 7** is not necessary. However, if the band is being analyzed in segments, then the tuner is used to step through a sequence of analysis segments of the full band, and the individual spectrum segments are spliced together to produce the overall wideband spectrum.

4.2. FBC Controls

The cable modem FBC feature is defined by the CmSpectrumAnalysis management information base (MIB) object in the CableLabs DOCSIS 3.1 CM OSSS specification [Ref. 3]. The CmSpectrumAnalysisCtrlCmd MIB provides a set of attributes which are used to configure and enable a spectrum capture. The cable modem spectrum capture is performed over a set of spectral segments where the segment width is defined by the SegmentFrequencySpan and the entire capture is defined by the FirstSegmentCenterFrequency, LastSegmentCenterFrequency and NumBinsPerSegment. WindowFunction allows selection from an assortment of windows. The RBW is a function of the NumBinsPerSegment but is scaled by an amount related to the WindowFunction in use for the capture; the reported RBW (units of Hz) is used to multiply an input signal or noise flat power spectral density (PSD, or dBmV per Hz) to yield the power (dBmV) measured in each bin: Flat PSD (dBmV/Hz) times RBW (Hz) equals bin power (dBmV). The frequency-resolving capability of the window, which is similar to the frequency span of its passband, is different than the RBW, and is described by the EquivalentNoiseBandwidth MIB object. This object is a unitless number relating the ratio of the window

frequency span divided by the FFT bin spacing. At least one vendor has scaled the application of the window function such that the RBW is always equal to the bin spacing, but the frequency-resolving capability is still described by the EquivalentNoiseBandwidth MIB object. An averaging feature, NumberOfAverages, is available for the FBC which uses the leaky integrator method. The averaging feature is very useful in that the averaging is performed by the cable modem software prior to making the data available. The data for a particular capture can be retrieved using the CmSpectrumAnalysisMeas object or by using the Bulk Data CM Spectrum Analysis File uploaded via TFTP. Additional details are discussed in Section 6.

The CableLabs specification has no requirements with respect to the accuracy of the spectrum analysis feature. However, the measurement inherits accuracy from the cable modem requirement to report the power per 6 MHz in any downstream channel with an accuracy of ± 3 dB. The amplitude data which is made available as just described uses 16 bit two's complement notation in units of hundredths dBmV per bin.

5. FBC Examples

This section includes several examples of FBC screen shots from operating cable networks. The examples highlight some of the more common impairments that can be identified using FBC. In systems where FBC is being used, technicians' comment that the spectrum analyzer-like displays are very intuitive and easy to interpret. Use case examples with more detailed information about the problems observed are included, too.




Figure 8. Adjacency – While the spectrum is relatively flat, note that signals below 400 MHz are a few dB higher in amplitude than signals above 400 MHz. This suggests incorrect RF level adjustment in the headend or hub site (courtesy of Comcast).




Figure 9. Filter – Indicates the presence of a filter in the subscriber drop (courtesy of Comcast).




Figure 10. Negative tilt – Typical response at or near ends-of-line locations, but excessive negative tilt could indicate a problem. Note presence of what appears to be FM broadcast band ingress at the left end of the display (courtesy of Akleza).




Figure 11. Positive tilt – Typical response at or near the output of nodes and amplifiers. Note presence of what appears to be FM broadcast band ingress at the left end of the display (courtesy of Akleza).




Figure 12. Standing wave – Classic example of scalloped sinusoidal wave shape in the response, caused by an impedance mismatch (courtesy of Akleza).




Figure 13. Suckout – Notch in the response, centered just above 700 MHz (courtesy of Akleza).




Figure 14. Suckout - A severe example, centered between 500 MHz and 600 MHz (courtesy of Comcast).




Figure 15. Water damage - Non-periodic wave shape in the response and higher attenuation at higher frequencies, typical of water damage in a subscriber drop (courtesy of Akleza).

The next four figures show resonant peaking varying over a four-day period. An important note: Should a varying response problem such as resonant peaking occur in the vicinity of the cable network's AGC pilot, amplifier operation will be affected and signal levels will vary, too. Tracking down what is often an intermittent problem like this is usually difficult, but FBC can be a valuable tool to help quickly identify and locate the source.




Figure 16. Resonant peaking – Day 1 FBC spectrum (courtesy of Comcast).




Figure 17. Resonant peaking - Day 2 FBC spectrum, same modem as previous figure (courtesy of Comcast).




Figure 18. Resonant peaking – Day 3 FBC spectrum, same modem as previous two figures. Note that the problem seems to have disappeared (courtesy of Comcast).




Figure 19. Resonant peaking - Day 4 FBC spectrum, same modem as previous three figures. The response problem has returned (courtesy of Comcast).

5.1. FBC Use Cases

The following examples include additional details related to the use of FBC to troubleshoot problems.

5.1.1. Water damaged drop

FBC was used to identify an individual subscriber drop that had water damage, specifically by the unique signature in the displayed frequency response. As **Figure 20** shows, the response has a non-periodic wave shape, and attenuation increases dramatically at higher frequencies. Remote polling of the modem showed that most of the downstream SC-QAM signals had poor performance, as shown in **Figure 21**. The upstream was relatively unaffected.




Figure 20. Water damage before repair. Note the non-periodic wave shape in the FBC response and the higher attenuation at higher frequencies. This particular example occurred when abrasion damaged the cable's jacket, allowing water to enter the cable (see text). Courtesy of Comcast.

Device Health																								
Registration State	6 (Online)																							
Down Rx Power	-18.5	-24.7	-28.3	-19.5	-19.3	-14.7	-14.9	-17.9	-22.2	-26.1	-26.1	-22.8	-19	-14.7	-12.5	-14.5	-16	-15.2	-17.7	-21.4	-21.5	-20.5	-21.5	-21.7
Downstream SNR	31.9	25.5	24	30.3	30.9	34.9	34.9	32.3	28.3	25.5	25.5	28.4	31.1	35.5	35.5	34.3	33.3	34.9	32.9	26.7	29.8	30.6	29.7	29.7
Upstream Tx Power	43.9						46.5						46						43.5					
Upstream SNR CM	34.6						33.6						36.2						36.6					
Upstream Rx Power	0						0						0						0					
US RX/NO Padding	0						0						0						0					
Upstream SNR Ch	34.6						33.7						36.4						36.7					
Upstream Ranging	4 (Success)						4 (Success)						4 (Success)						4 (Success)					

Figure 21. Most of the downstream SC-QAM signals have low signal level and degraded RxMER, indicated in red shaded boxes. The upstream was relatively unaffected (courtesy of Comcast).

Technicians went to the subscriber location and were able to find and fix the problem without having to enter the premises. **Figure 22** shows the coax jacket, which had been damaged by abrasion from the electrical service drop. The damaged coax jacket allowed water to enter the cable and travel inside of the cable all the way to the ground block. **Figure 23** shows water coming out of the connector at the ground block end of the drop.



Figure 22. Damaged coax jacket where water was able to enter the cable.



Figure 23. Water coming out of the end of the connector at the ground block.

The fix was to replace the subscriber drop from the tap to the ground block. **Figure 24** shows the FBC screen shot after the new drop was installed, and **Figure 25** the post-repair SC-QAM performance.




Figure 24. FBC response after drop cable was replaced from the tap to the ground block (courtesy of Comcast).

Device Health																								
Registration State	4 (Online)																							
Down Rx Power	1	1	0.9	1.2	1	1	1	0.9	0.7	0.9	0.5	0.9	1	1	1	0.9	0.7	0.9	1	1.2	1	0.7	0.7	0.4
Downstream SNR	40.3	40.9	40.9	40.8	40.3	40.9	40.3	40.8	40.3	40.9	40.3	40.3	40.9	40.9	40.9	40.3	40.9	40.9	40.9	40.9	40.3	40.3	40.9	40.3
Upstream Tx Power	39.8						40.3						40.5						40.3					
Upstream SNR CM	34.8						35.7						36.5						36.6					
Upstream Rx Power	0						-0.1						-0.1						-0.2					
US RX/WO Padding	0						-0.1						-0.1						-0.2					
Upstream SNR Ch	34.8						35.7						36.5						36.6					
Upstream Ranging	4 (Success)						4 (Success)						4 (Success)						4 (Success)					

Figure 25. Signal performance after the drop was replaced (courtesy of Comcast).

Additional discussion about water-soaked cable can be found in Section 7.4.

5.1.2. Localizing a problem to a specific drop

A problem was detected at Subscriber A: The OFDM modulation profile was limited to 64-QAM, but was expected to be 2048-QAM or better.

The OFDM signal in this 860 MHz HFC network is from 774 MHz to 864 MHz, with an exclusion band configured – 790 MHz to 820 MHz – to avoid ingress from the European 800 MHz 4G/LTE band.

Using the cable modem's built-in FBC function (port 8080), strong ingress was detected at all 5G and LTE frequencies in the area, as shown in **Figure 26**.




Figure 26. Subscriber A's FBC showing 4G/LTE and 5G ingress (courtesy of Telia Norge).

Looking up the neighboring address (Subscriber B) and using FBC in that cable modem (**Figure 27**), no ingress is to be seen.




Figure 27. Subscriber B's FBC, showing no 4G/LTE or 5G ingress (courtesy of Telia Norge).

It can now easily be determined that the root cause of the problem is related to the drop cable and/or house internal network at Subscriber A's premises.

Checking the location of nearby base stations (**Figure 28**), it turns out that Subscriber A is located approximately 420 meters (~1380 feet) from a cell site that is operating on the 700, 800 and 900 MHz bands.




Figure 28. Location of cell site near the affected subscriber drop.

5.1.3. FBC Analysis and Fault Location Example

In this example a standing wave was detected that impacting a large part of a node's service area. This problem could clearly be seen in the FBC plots as shown in the traces on the right-hand side of **Figure 29** and was impacting all subscribers downstream of amplifier [1]. Modems that were being fed by a different distribution leg of upstream amplifier [2] did not show the impairment as can be seen in the FBC plot on the left. This analysis immediately isolates the problem to output connections coming from amplifier [1], the trunk output of amplifier [2] connecting to amplifier [1], or a cable fault between the two amplifiers.

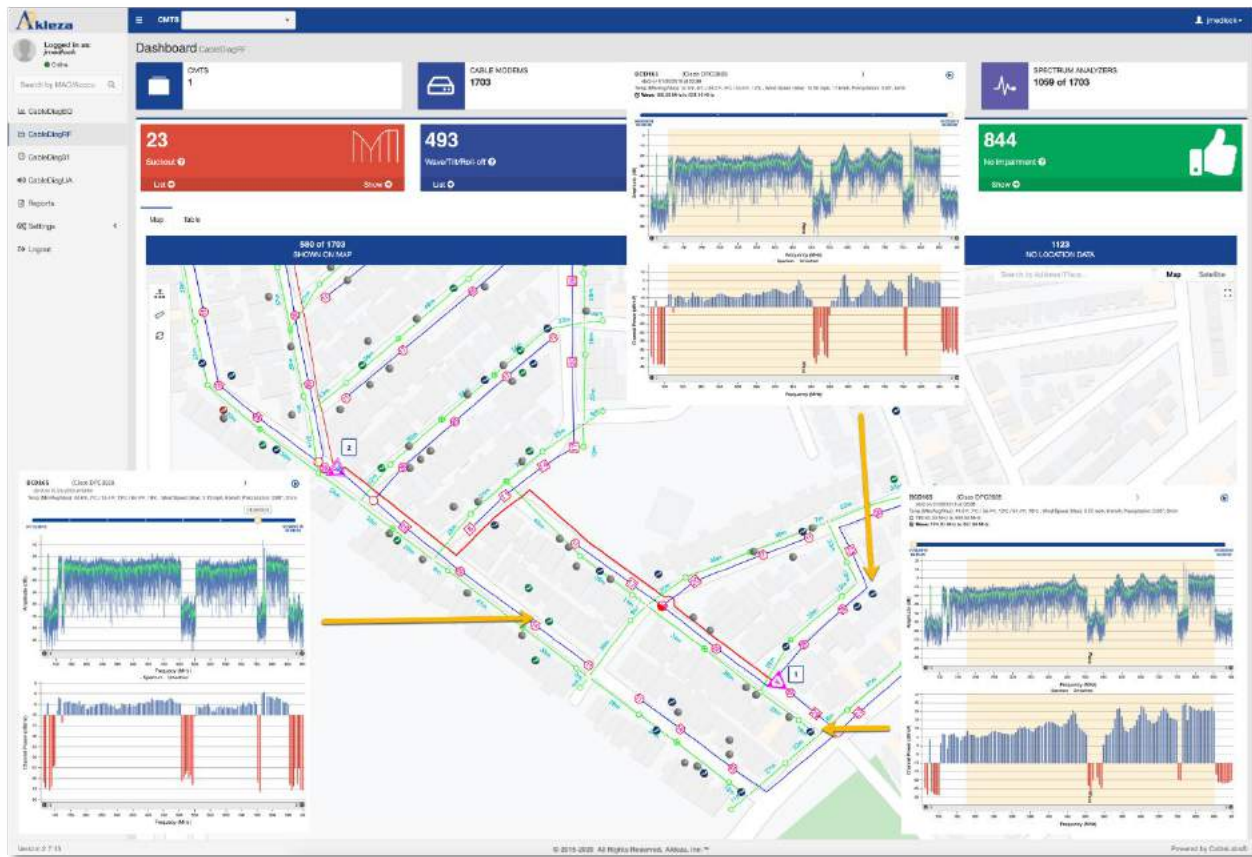


Figure 29. Example of using FBC to isolate a standing wave fault (courtesy of Akleza).

The technicians first checked amplifier [1] and its output connections. When the amplifier was opened, water poured out indicating a problem with the enclosure seal. Additionally, the internals of the amplifier showed corrosion as can be seen in **Figure 30**.



Figure 30. Amplifier [1] containing water and showing corrosion (courtesy of Akleza).

Given the amount of damage to the amplifier, it was replaced, and the cable modems rescanned. This did not correct the problem indicating that the fault causing the standing wave was upstream of amplifier [1]. A field meter attached to the input side of the amplifier was used to verify this analysis and showed the problem could be seen in the signal coming in from upstream amplifier [2] (**Figure 31**).

The fault must therefore be with upstream amplifier [2] or its trunk output as it was already determined that the fault could not be seen on either of the amplifier's auxilliary distribution outputs.




Figure 31. Fault still visible after replacement of amplifier [1] (courtesy of Akleza).

Verifying the signal at the output test port for the trunk connection did not show any standing wave meaning that the problem had to be with the output connector itself or the cable. When the output connection was inspected it was noticed that the connector showed signs of corrosion as well as the center conductor being cut too short, and therefore not making good contact internally (**Figure 32**). Corrosion and badly installed connectors are common causes of impedance mismatches in the cable plant and can cause standing wave spectrum faults.



Figure 32. Amplifier [2] with corroded connector (courtesy of Akleza).

Once the output connector was replaced, the cable modems were rescanned and showed that the standing wave fault had been fixed. **Figure 33** shows the resulting FBC plot from one of the cable modems after the fault was corrected.

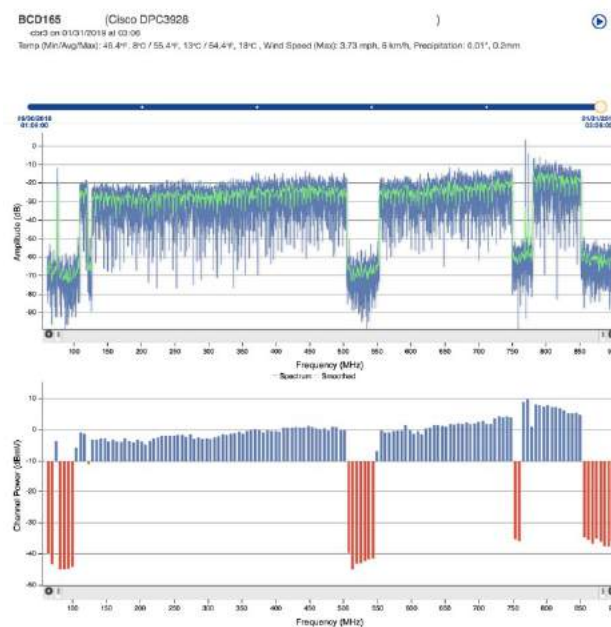


Figure 33. FBC trace after faulty connector replaced (courtesy of Akleza).

This example shows how analyzing FBC data collected remotely from cable modems within the network can quickly be used to identify customer service faults and quickly triangulate the fault location based on the FBC response from multiple devices.

6. How to retrieve and Display FBC Data

6.1. FBC Data Collection

In DOCSIS 3.0 the configuration and collection of FBC data uses SNMP requests to a cable modem supporting FBC. The PNM Best Practices: HFC Networks (DOCSIS 3.0) guide [Ref. 4] describes the high-level functionality while the SNMP MIB objects are defined in the DOCS-IF3-MIB.³

DOCSIS 3.1 introduced a new set of requirements on the cable modem related to the collection and reporting of PNM data as described in Section 9 of the DOCSIS 3.1 Physical Layer Specification [Ref. 5], and in PNM Best Practices Primer: HFC Networks (DOCSIS® 3.1) [Ref. 8]. DOCSIS 3.1 also introduced a new bulk-data transfer mechanism that uses the trivial file transfer protocol (TFTP) to upload PNM data to a destination TFTP server. Configuration of the PNM test execution, file storage, and TFTP destination use SNMP as defined in [Ref. 3] and CCAP Operations Support System Interface (OSSI) Specifications [Ref. 6].

While the DOCSIS 3.1 bulk-data transfer mechanism supports the collection of FBC data, this is only supported by DOCSIS 3.1 cable modems. The DOCSIS 3.0 SNMP-only mechanism is, however, supported by both DOCSIS 3.0 and 3.1 modems and functionally provides the same capability so can be used as a single method across both cable modem types. This section covers the use of the DOCSIS 3.0 SNMP configuration and retrieval mechanism. For more information on the DOCSIS 3.1 bulk-data transfer mechanism refer to [Ref. 3] and the DOCS-PNM-MIB.

Commercial FBC data collection, analysis, and display products are generally available in the market from PNM vendors. The following information is provided on how to configure, retrieve, and plot FBC data for those looking to experiment with and understand what is possible with that data. Because capturing spectrum information requires SNMP read/write access to the cable modems, access using a locally attached workstation is unlikely. Typically, an FBC solution is deployed within the cable operator's operations or back office organization and must deal with a number of elements including security and data storage. A high resolution, full downstream spectrum capture can generate 200 kB to 300 kB of data that is streamed back from the cable modem. Deploying a system-wide solution therefore requires careful engineering as well as suitable resources to handle the data collection and storage requirements.

In the provided examples the freely available Net-SNMP tools are used which can be installed on most operating systems. For specific instructions on installing and using Net-SNMP for a particular platform visit the Net-SNMP website at www.net-snmp.org. Any SNMP community string or IP addresses included in this paper are simply examples and should be updated to match your own environment. Some Python-based scripts have also been provided to assist in the collection and display of the FBC data. These are provided purely for informational purposes and should not be considered production-level code. **Figure 34** shows the DOCS-IF3-MIB objects that will be used to configure and retrieve the FBC data from the cable modem.

³ The CableLabs MIB repository is here: <http://mibs.cablelabs.com/MIBs/DOCSIS/>

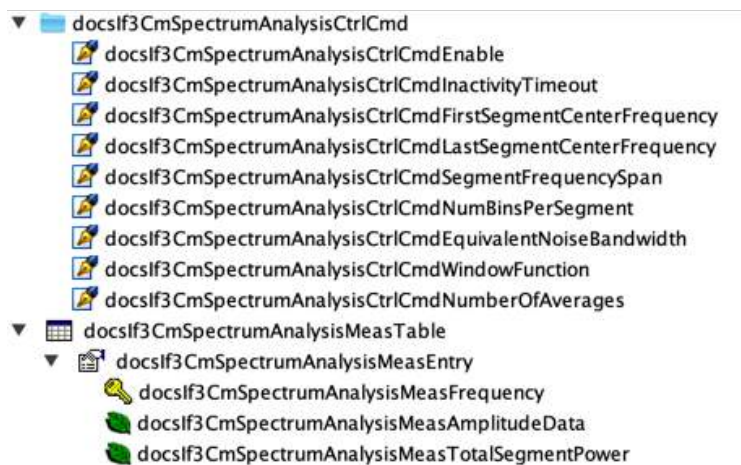


Figure 34. IF3-MIB spectrum capture objects

While in general all DOCSIS 3.0 and 3.1 cable modems should be capable of spectrum capture, some early DOCSIS 3.0 cable modems running older versions of firmware may not support the capability. In order to first determine if a cable modem supports FBC, a simple SNMP read operation on the docsIf3CmSpectrumAnalysisCtrlCmdEnable object can be performed. If the device supports FBC, this object will return a value, while a device that does not support the function will return an SNMP No Such Object error meaning that the cable modem firmware does not support the DOCS-IF3-MIB FBC objects.

This example shows a response from a cable modem that supports FBC.

```
$ snmpget -v2c -c public 10.60.0.7 docsIf3CmSpectrumAnalysisCtrlCmdEnable.0
DOCS-IF3-MIB::docsIf3CmSpectrumAnalysisCtrlCmdEnable.0 = INTEGER: false(2)
```

This example shows the response from a cable modem that does not support FBC.

```
$ snmpget -v2c -c public 10.60.0.17 docsIf3CmSpectrumAnalysisCtrlCmdEnable.0
DOCS-IF3-MIB::docsIf3CmSpectrumAnalysisCtrlCmdEnable.0 = No Such Object available on this agent at this OID
```

Instructing the cable modem to collect FBC data and then retrieving the data is a two-step process. First the docsIf3CmSpectrumAnalysisCtrlCmd objects need to be set to values describing the data capture you want, and then the results are read from the docsIf3CmSpectrumAnalysisMeasTable. Again, depending upon the vintage of the cable modem and the functionality supported by its firmware, not all configurations may be possible.

To understand the configuration of the spectrum capture, you first need to consider how much of the spectrum you want to capture, and the resolution of the captured data. Between the start (docsIf3CmSpectrumAnalysisCtrlCmdFirstSegmentCenterFrequency) and end (docsIf3CmSpectrumAnalysisCtrlCmdLastSegmentCenterFrequency) points of the desired spectrum capture, you define the width of a segment (docsIf3CmSpectrumAnalysisCtrlCmdSegmentFrequencySpan) and then how many samples or bins (docsIf3CmSpectrumAnalysisCtrlCmdNumBinsPerSegment) you want measured within that segment. The configuration of these parameter defines how much data is going to be collected and needs to be read back from the cable modem. The greater the number of bins the higher the spectral resolution, however, this will generate more data that must be retrieved.

For example, if you request a capture of the full downstream spectrum, from, say, 54 MHz to 1002 MHz, using a segment frequency span of 6 MHz, and 256 bins per segment, you would generate 40,192 data points at a resolution of 23 kHz per point. Reduce the number of bins to 64, and you would now only

have 10,048 data points with a resolution of 94 kHz per point. Depending upon how you intend to use the data and the resources available to capture and store the data, you can adjust these configuration parameters accordingly.

To initiate a spectrum capture, you first must set the configuration parameters and the enable flag. **Figure 35** shows how to configure a high resolution full downstream data collection as just described.

```
$ snmpset -v2c -c public 10.30.0.12 \
docsIf3CmSpectrumAnalysisCtrlCmdFirstSegmentCenterFrequency.0 u 57000000 \
docsIf3CmSpectrumAnalysisCtrlCmdLastSegmentCenterFrequency.0 u 999000000 \
docsIf3CmSpectrumAnalysisCtrlCmdSegmentFrequencySpan.0 u 6000000 \
docsIf3CmSpectrumAnalysisCtrlCmdNumBinsPerSegment.0 u 256 \
docsIf3CmSpectrumAnalysisCtrlCmdEnable.0 i 1
```

Figure 35. Configuration example for high resolution downstream data collection.

The command in **Figure 35** sets the span width to 6 MHz, the start center frequency to 57 MHz, and the end center frequency to 999 MHz. Setting the enable value to true (1) signals the cable modem to start its spectrum capture. Using the 6 MHz span width additionally allows the spans to be aligned to standard channel frequencies whose power is also measured during the spectrum capture.

With the spectrum capture now enabled, the spectrum data can be retrieved by doing a snmpwalk of the docsIf3CmSpectrumAnalysisMeasAmplitudeData. This will return an SNMP table row for each frequency span as shown in **Figure 36**.

```
$ snmpwalk -v2c -c public 10.30.0.12 docsIf3CmSpectrumAnalysisMeasAmplitudeData
DOCS-IF3-MIB::docsIf3CmSpectrumAnalysisMeasAmplitudeData.57000000 = Hex-STRING:
03 65 C0 40 00 5B 8D 80 00 00 01 00 00 00 5B 8D
00 00 5B 8D E6 D6 E6 EA E6 9D E6 43 E6 11 E6 19
E5 A6 E4 56 E3 B8 E4 59 E5 B8 E6 41 E6 5C E6 2B
: SOME DATA ROWS REMOVED TO SAVE SPACE
E6 C7 E6 B0 E6 7B E6 4B E6 06 E5 73 E4 8D E3 77
E0 F5 E0 8E E1 AC E3 41 E3 D1 E3 92 E3 4C E2 52
E1 47 E2 1C
:
DOCS-IF3-MIB::docsIf3CmSpectrumAnalysisMeasAmplitudeData.999000000 = Hex-STRING:
3B 8B 87 C0 00 5B 8D 80 00 00 01 00 00 00 5B 8D
00 00 5B 8D E6 28 E4 0A E1 FB E2 9D E5 59 E6 59
E6 86 E6 AB E6 8E E5 FC E4 C9 E2 EE DE 32 DE 5F
: SOME DATA ROWS REMOVED TO SAVE SPACE
E8 87 E8 8F E8 26 E7 B9 E7 2F E6 E3 E6 C5 E6 61
E5 3E E4 55 E5 31 E7 45 E7 FC E8 2F E7 D3 E7 0C
E5 7A E2 41
```

Figure 36. Returned SNMP table row for each frequency span.

The first 20 bytes of each rows data consists of five integers (4 bytes) as follows:

- 4 bytes: CenterFreq
The center frequency of the span.
- 4 bytes: FreqSpan
The width in Hz of the span.
- 4 bytes: NumberOfBins
The number of data points or bins that compose the spectral data. The leftmost bin corresponds to the lower band edge, the rightmost bin corresponds to the upper band edge, and the middle bin center is aligned with the center frequency of the analysis span.
- 4 bytes: BinSpacing
The frequency separation between adjacent bin

centers. It is derived from the frequency span and the number of bins or data points. The bin spacing is computed as:

$$\text{BinSpacing} = \text{FrequencySpan} / (\text{NumberOfBins} - 1)$$

The larger the number of bins the finer the resolution.
4 bytes: ResolutionBW
The resolution bandwidth or equivalent noise bandwidth of each bin. If spectral windowing is used (based on vendor implementation), the bin spacing and resolution bandwidth would not generally be the same.

Each remaining two-byte data pair represents a bin amplitude in units of 0.01 dB. The bin amplitude data is in two's complement format so must be translated into a decimal value. There are a number of ways to convert two's complement data to decimal, but to simplify, we first convert from hexadecimal to decimal. Then, if the value is greater than 32767, it represents a negative number so we subtract 65535; otherwise, use the initially converted value. Divide the result by 100 to get the amplitude bin in dB. For example:

E6D6 in decimal is 59094
As $59094 > 32767$ then bin value = $59094 - 65535 = -6441$
Bin Amplitude in dB = $-6441 / 100 = -64.41$ dB

Using the example data above, the first row would decode as:

center frequency: 0365C040 = 57000000 Hz
frequency span: 005B8D80 = 6000000 Hz
number of bins: 00000100 = 256
bin spacing: 00005B8D: 23437 Hz
resolution bandwidth: 00005B8D = 23437 Hz
Bin 0: E6D6 = -64.41 dB
Bin 1: E6EA = -64.21 dB
Bin 2: E69D = -64.98 dB
:
Bin 255: EC1C = -50.91 dB

Given the center frequency, the number of bins, and the bin spacing, the frequency of each bin can be calculated

$$\text{Frequency} = \text{center frequency} - (\text{number of bins} / 2 * \text{bin spacing}) + (\text{bin number} * \text{bin spacing})$$

Using this formula, the frequency for the above example bins would be:

Bin 0: 54000064 Hz
Bin 1: 54023501 Hz
Bin 2: 54046938 Hz
:
Bin 255: 59976499 Hz

After decoding each of the `docsIf3CmSpectrumAnalysisMeasAmplitudeData` rows into their frequency and bin amplitude parts, you now have a list of frequencies and amplitudes that can be plotted to show the FBC spectrum.

In order to simplify the retrieval, decoding, the Appendix contains Python code listings of scripts that can be used to collect and display FBC data. The output of the `getFbcData.py` script is a comma separated values (CSV) file containing frequency and amplitude columns. This file can be directly imported into Excel and a chart created. **Figure 37** shows an example Microsoft Excel chart using this imported data.




Figure 37. FBC plot using Microsoft Excel (courtesy of Akleza).

The Python script `showFbcData.py` reads a previously saved FBC CSV file and displays a plot using the Python Plotly package. **Figure 38** shows the same FBC data displayed using this script.




Figure 38. FBC plot generated using example Python scripts (courtesy of Akleza).

As shown in **Figure 34** there are a number of configuration parameters that can be used to control the extents and resolution of the FBC function on the cable modem. Configuration parameters are also

available to control the spectral windowing function and the number of captures to be averaged during the spectrum capture.

The DOCS-PNM-MIB defines a number of windowing filters to smooth out spectral leakage problems in the underlying capture mechanism. While a number of window functions are specified, not all cable modem implementations support all of them but most do support the Hanning window which is commonly used by setting the `docsIf3CmSpectrumAnalysisCtrlCmdWindowFunction` to `hann(1)`.

Averaging is an important configuration option to help in smoothing the instantaneous capture of the underlying ADC within the FBC process. For example, **Figure 39** shows the result of a capture with no averaging performed.

```
$ snmpset -v2c -c public 10.30.0.8 \  
docsIf3CmSpectrumAnalysisCtrlCmdFirstSegmentCenterFrequency.0 u 57000000 \  
docsIf3CmSpectrumAnalysisCtrlCmdLastSegmentCenterFrequency.0 u 999000000 \  
docsIf3CmSpectrumAnalysisCtrlCmdSegmentFrequencySpan.0 u 6000000 \  
docsIf3CmSpectrumAnalysisCtrlCmdNumBinsPerSegment.0 u 256 \  
docsIf3CmSpectrumAnalysisCtrlCmdWindowFunction.0 i 1 \  
docsIf3CmSpectrumAnalysisCtrlCmdNumberOfAverages.0 u 1 \  
docsIf3CmSpectrumAnalysisCtrlCmdEnable.0 i 1
```




Figure 39. FBC plot with no averaging (courtesy of Akleza).

Figure 40 shows a capture where the `docsIf3CmSpectrumAnalysisCtrlCmdNumberOfAverages` was set to 32. Now you can more clearly see details such as the pilots and PLC in the OFDM channel which are lost in **Figure 39**.

```
$ snmpset -v2c -c public 10.30.0.8 \
docsIf3CmSpectrumAnalysisCtrlCmdFirstSegmentCenterFrequency.0 u 57000000 \
docsIf3CmSpectrumAnalysisCtrlCmdLastSegmentCenterFrequency.0 u 999000000 \
docsIf3CmSpectrumAnalysisCtrlCmdSegmentFrequencySpan.0 u 6000000 \
docsIf3CmSpectrumAnalysisCtrlCmdNumBinsPerSegment.0 u 256 \
docsIf3CmSpectrumAnalysisCtrlCmdWindowFunction.0 i 1 \
docsIf3CmSpectrumAnalysisCtrlCmdNumberOfAverages.0 u 32 \
docsIf3CmSpectrumAnalysisCtrlCmdEnable.0 i 1
```




Figure 40. FBC plot with 32 averages (courtesy of Akleza).

Unfortunately, older cable modems do not fully support averaging, some supporting a smaller number of averages, and some none at all. In these cases, the initial set of the `docsIf3CmSpectrumAnalysisCtrlCmdNumberOfAverages` object will fail with an SNMP error as follows:

```
Reason: inconsistentValue (The set value is illegal or unsupported in some way)
Failed object: DOCS-IF3-MIB::docsIf3CmSpectrumAnalysisCtrlCmdNumberOfAverages.0
```

When this occurs the polling application will have to make multiple requests to retrieve a number of samples and then implement a running average over these samples.

With both averaging and collection of high-resolution captures, significant detail can be seen. **Figure 41** shows the same capture as above but zoomed into see the detail of the SC-QAM signals and the start of the OFDM signal with its pilots and PLC.




Figure 41. Zoomed view of high resolution and averaged FBC (courtesy of Akleza).

7. Advanced FBC Applications

7.1. Estimating Distances with High Accuracy

Because of a large potential bandwidth associated with an FBC spectrum, the ripples or waves in the magnitude plot can be analyzed to produce very accurate distance estimates. A signal captured in a narrow band can only produce inaccurate distance measurements. One method that has been developed in the CableLabs PNM Working Group is to process spectral data with an inverse discrete Fourier transform, using the FBC response for real coefficients and all zeroes for the imaginary coefficients. If a broadband signal is transformed, this method produces a highly accurate time response, and a magnitude response with decent and predictable results.

It is useful to examine exactly how delayed signals arrive at the cable modem or time domain reflectometer (TDR), and how and where they can be observed using, for example, high impedance probes. **Figure 42** has two diagrams showing how reflections are generated on a cable line, one with a single reflection and one with two reflections. The cable lines are blue and X marks the locations of impedance mismatches R, RA, and RB. Each impedance mismatch in this example causes a 20% reflection ($S_{11} = S_{22} = 0.2$). The horizontal lines are signals going downstream to the right and diagonal lines are signals going upstream to the left. The strength of the signals in volts for each cable segment is labeled in red and the coax is assumed to have no cable loss. The TDR impulse response plot is shown on the left, and an FBC impulse response received downstream is shown on the right. Upstream and downstream responses can also be observed by high impedance probes, placed as illustrated.

For the single reflection case, a 20% signal is reflected back from point R and is observed at probe 1, but no delayed copy of the signal appears downstream (probe 2). However, probe 2 does observe that the downstream signal is weaker than it should have been (0.8 instead of 1.0). In other words, FBC on a cable modem is effectively blind to a single reflection caused by a single impedance mismatch.

For the two-reflection case, signals bounce back and forth between reflection points RA and RB infinitely, although generally only the first few recursions can be observed before the signal is too weak to be seen. A ripple in the frequency domain is equivalent to one or more impulses in the time domain. Thus, when a ripple in the frequency domain is observed, you know you have an echo tunnel (distance between RA and RB), but you don't know the locations of the start point or stop point. However, you can determine an exact length from an exact time delay obtained by an inverse Fourier transform. **Figure 43** shows water damage that created one end of an echo tunnel, and the other end of the tunnel was caused by a seizure screw not being tightened on a terminator.

Additional information about reflection types can be found in [Ref. 7].




Figure 42. Diagrams showing what is observed at two points on a cable line for the one- and two-reflection cases. The impulse responses on the right can be observed by the FBC and the impulse responses can be observed by an active or passive TDR.



Figure 43. Damaged hardline coax that was at one end of an echo tunnel; the other end was a loose seizure screw on a chassis terminator installed in an end-of-line tap. Further

complicating things: The damaged cable was submersed in water. Courtesy of CableLabs.

7.2. Drop Cable Testing

A simple test that can be done on drop cable to verify that it is working well is to disconnect one end, say at the ground block, and test the drop cable from the tap with a TDR. The TDR impulse response should reveal the open drop cable with only a few dB of attenuation associated with the drop cable. Any other response indicates a bad drop cable. This method can be performed with a conventional metallic TDR or with a passive TDR.⁴

7.3. Using FBC to View Upstream Noise

The FBC's tuning range includes the upstream band. The upstream band is of interest to capture for much the same reasons as the downstream band; just as impairments add energy that interferes with reception of downstream signals (see Section 3), energy can be added that interferes with the reception of upstream signals. Assume that somewhere in the drop or house wiring is a shield break that allows electrical noise traveling on the outside of the coax to gain entry to the coax wiring. Once the noise gains entry, it travels in both directions – toward the CPE and toward the headend or hub – the latter where it can cause interference to other upstream signals at the CMTS receiver. The same interference can be detected by terminal equipment located in the home, identifying the home where the ingress entered the cable plant.

How strongly the ingress enters the plant, compared to the upstream transmission, will determine the ingress-to-signal ratio at the upstream receiver. Also, where the ingress enters the plant, in terms of insertion loss to-and-from the cable modem, will determine how the ingress-to-signal ratio will appear at the F-connector of the cable modem. Even if there are no filters or active components in the path between the ingress point and the cable modem, the ingress-to-signal ratio will be lower at the cable modem F-connector than at the CMTS upstream receiver, and will be lower by at least twice the insertion loss of the path from the cable modem to the ingress point.

An additional consideration is that upstream signals can have levels in the range of +17 dBmV/1.6 MHz to +53 dBmV/1.6 MHz, generally; for comparison, downstream signal levels can be in the range of -10 dBmV to +15 dBmV per 6 MHz channel. In some cases the cable modem's transmitted upstream energy can be significantly greater than it is receiving in the downstream. To be useful in locating ingress it will generally be necessary to capture upstream ingress when upstream transmissions are not occurring at the same frequency. This is because often an impairing ingress has smaller PSD than the transmitted signal, at the cable modem, and will be obscured by the cable modem's own transmission if it is occurring at the same frequency as the ingress.

Capturing the upstream spectrum at the F-connector of the cable modem can help provide insight into troubleshooting, and lead to fixing the cause of the ingress. Since the ingress-to-signal ratio is generally already lower at the cable modem than at the upstream receiver, it is desired to capture the upstream spectrum without the stopband loss that would occur on the high-frequency port of a cable modem's

⁴ Water damaged cable may not show up well unless the TDR has test energy in the UHF spectrum, which can usually be achieved by choosing a very narrow pulse width.

internal duplex filter. Details of how this is facilitated in some cable modem implementations are beyond the scope of this paper.

7.4. Using FBC to Find Water-Soaked Coaxial Cable

Figure 44 is a frequency response of a coaxial cable that has water inside the jacket. It can be compared to **Figure 45** which is caused by a standing wave created by a pair of separated impedance mismatches. Operational experience has shown that this water characteristic is a highly reliable indicator of the presence of water in coax, most typically in drop cable. But water responses have also been observed in hardline coax. If a coaxial cable is tested with a TDR in wide bandwidth mode (narrow pulse width), a nearby region of water damage can be localized. It has been observed to correct itself when the coax either dries out or the water inside is frozen hard, although its characteristic reappears when the frozen water thaws. Water gets inside the coax at abrasion points, hairline cracks, kinks, or at connectors. The rough non-periodic spectral response is caused by the non-uniform saturation of water-versus-distance over a section of the coaxial cable. Generally, a second characteristic accompanies the rough spectral response: higher attenuation at high downstream frequencies, as much as 15 dB to 25 dB (or more!). Upstream attenuation is typically only a few dB, but the high downstream attenuation often results in greatly reduced throughput and customer dissatisfaction. The adaptive equalizer in the cable modem can tolerate the rough spectral response, but loss of signal level results in data errors.

Relatively simple software currently exists to detect both types of impairments by their rapid variation of amplitude versus frequency, but these simple algorithms cannot distinguish water damage from standing waves.

A DSP algorithm is being developed where a spectral response is filtered, flattened and interpolated to remove high-level responses such as AGC pilot signals, and low responses such as vacant spectrum. Finally, an inverse Fourier transform produces an impulse response which is examined for dispersion. One or two narrow lines indicate a standing wave, and a dispersed time response indicates water. The dispersion can be enhanced by performing an autocorrelation function on the time response, improving detection. Another revealing characteristic is large spectral down-tilt caused by water attenuation. **Figure 44** shows a response with water in time and frequency and **Figure 45** shows the same data for a standing wave.

One simple method to measure dispersion is to measure the impulse response coefficients relative to the DC term, remove the two largest coefficients, and then remeasure the coefficients. A small drop after coefficient removal indicates a dispersed water response and a large drop indicates a large standing wave response. Another method is to take the ratio of the peak-to-average of the time coefficients, where a large peak-to-average indicates a standing wave, and a smaller peak-to-average ratio indicates a non-periodic water-caused wave. Note that it is possible for a cable to suffer both water and standing wave impairments at the same time. Likewise, it is possible for a cable to have a second standing wave frequency present.




Figure 44. Water-soaked coax, FBC plot on left and impulse response on right. The impulse response is highly dispersed.




Figure 45. Coax with standing wave response. FBC plot is on left and impulse response is on right. The impulse response is not dispersed.

8. Impairment Detection

Detecting impairments from spectrum analysis data is simple but time consuming to do manually. Fortunately, there are methods that can assist with this problem. Three practical categories for addressing this problem effectively include:

- Simple magnitude and statistical methods followed up with manual inspection,
- Mathematical and statistical models to recognize impairments followed up with manual inspection, and
- Machine learning approaches that find and classify impairments, for either direct action or followed up with manual inspection.

Only operators will have numbers to compare between these methods, and so far we have not found any data being shared. We expect, however, that the first two methods will generally be slower than the third, based on a comparison of methods developed by CableLabs; and we expect that false positives and false negatives will decrease moving down this list. Further, because the purpose is to address the impairments, eventually someone will manually validate the results, before or during dispatch.

However, it is important to mention that automating the first inspection of spectrum data can reduce the bulk of manual labor necessary. The speed of automated inspection may not be a critical factor, but rather the ability to do a reliable first inspection will save a large amount of manual work.

For example, consider 1000 results. If there are 10 impairments in 1000 samples, a manual process has to look at all 1000 samples to find them. An automated process might have 50% false positives and 10% false negatives, so might report 18 impairments. After looking at these 18 results manually, which is far faster than looking at 1000, nine impairments may be found. The net result is finding nine impairments by looking at 18 for a 50% efficiency, compared to finding 10 impairments by looking at 1000 for a 1% efficiency. While not all impairments are found due to the false negative rate, the efficiency is profound.

For some time, there have been DOCSIS® 3.0 anomaly detection methods available in the CableLabs Common Code Collection (C3). This method is a rule-based method that finds anomalies, identifies their types, and finds the frequencies they impact, all using rules, statistics, and traditional methods. CableLabs created a prototype which Comcast implemented and improved on, and that version is available for use by operators.

For DOCSIS 3.1, CableLabs offers ProOps as a platform for operators to test PNM methods, which includes a statistical method as well as the potential for a machine learning-based anomaly detection method. The latter is used for RxMER data today, but we are confident that it can be trained to work with spectrum data if operators wish to try this.

8.1. Impairment detection and automation

Automating FBC impairment analysis helps present the data in ways that are easy to understand and drive business decisions. Anomaly detection and classification algorithms are available through CableLabs or built into commercially available PNM tools. These algorithms detect and categorize the severity of the impairments at each FBC-capable cable modem in the network. In addition to the detection and severity categorization, the automation tools typically provide statistics and a list of modems with similar impairments grouped by geographic areas. This gives insight into the area of the HFC network that is of interest and has the potential to impact the customer experience of many subscribers.

For example, detected impairments are typically color coded by severity. Yellow for minor (less likely to be impacting) and red for major (more likely to be impacting). This enables the automation tool to assign a status of “red” to a modem that has one or more impairments identified, meeting the criteria for “red.” Subsequently it is easy to count and calculate the percentage of modems with a “red” status, either overall or by impairment. Combine the impairment severity information with HFC network topology information, and the automation tool will guide an engineer or network analyst to the HFC segment with the most issues identified.

The impairment algorithms use the FBC spectrum data, based on the bins, and calculate the power levels to represent the 6 MHz channels when required. The following is a brief description of the available impairment algorithms and their behavior:

Tilt – The condition where signal levels vary in a linear manner as the frequency increases. The variation can be low to high (positive tilt) or high to low (negative tilt); refer back to Figure 11 and Figure 10 respectively. The parameters of the detection can be adjusted, but typically tilt is identified when the power variation is calculated to be 15 dB or more. The algorithm will also calculate a severity number between 0 and 100, with 0 being the least severe and 100 being the most severe. In this way, the severity can be broken down into yellow or red. An example configuration would be yellow for severity number between 0 and 80, and red severity for 80 to 100.

Adjacency – Identifies a group of channels with median amplitude values several dB higher than other channels in that part of the spectrum, likely caused by incorrect narrowcast injection levels; refer back to Figure 8. The algorithm calculates the median amplitude within the 6 MHz channel and triggers an impairment detection when the adjacent median amplitude differs (typically by 3 dB to 14 dB). An example severity range would be yellow ≤ 50 and red > 50 .

Standing waves – Algorithm automation calculates changes in the median amplitude to determine when the signal includes amplitude ripple (standing waves); refer back to Figure 12. The algorithm can account for the natural tilt of the spectrum. Severity is calculated based on the difference between the min and max power level in the detected wave range. Example severity levels are yellow ≤ 7.5 dB and red > 7.5 dB.

Notch (suckout) – Algorithm detects when amplitude differences and high tilt are present in groups of channels; refer back to Figure 13 and Figure 14. Severity is based on the amplitude depth of the notch.

Rolloff – Can automatically detect non-flat loss of signal level at or near the upper end of the RF spectrum. The algorithm uses a rolling median and calculates severity by steepness of the roll-off amplitude over 4 dB.

Resonant peak – The algorithm detects areas of the spectrum where the amplitude maximum is higher than the average of the surrounding frequencies; refer back to Figures 16, 17, 18, and 19. Severity is determined by the difference between the maximum and the average amplitudes. Sample severity threshold is yellow at ≤ 7 dB and red > 7 dB.

9. Conclusion

Many DOCSIS 3.0 cable modems – and all DOCSIS 3.1 cable modems – support full band capture capability. FBC gives cable operators the equivalent of a spectrum analyzer in every home in which an FBC-capable modem has been installed. FBC spectral displays can show impairments such as ingress, frequency response problems, the presence of filters, incorrect narrowcast versus broadcast signal levels, and much more. Furthermore, those spectral displays are intuitive and easy to understand. A key benefit is that many problems can be identified remotely without the need to first roll a truck.

Some cable operators use FBC for plant maintenance and troubleshooting, and some operators and third-party vendors have incorporated FBC in their PNM tools. Yet many cable operators do not take advantage of FBC – technology that already exists in their networks – or are unaware of its capabilities.

The industry has been given a valuable and powerful tool for plant maintenance, and software exists to integrate FBC with operational practices. The opex is out there, waiting to be saved!

10. Abbreviations

4G	fourth generation [mobile telecommunications technology]
5G	fifth generation [mobile telecommunications technology]
ADC	analog-to-digital converter
AGC	automatic gain control
CCAP	converged cable access platform
CMTS	cable modem termination system
CPE	customer premises equipment
CSV	comma separated values

CW	continuous wave
dB	decibel
dBmV	decibel millivolt
DC	direct current
DFT	discrete Fourier transform
DOCSIS	Data-Over-Cable Service Interface Specifications
DSP	digital signal processing
FBC	full band capture
FFT	fast Fourier transform
FM	frequency modulation
HFC	hybrid fiber/coax
I	in-phase (real)
Hz	hertz
ISBE	International Society of Broadband Experts
kB	kilobyte
kHz	kilohertz
log	logarithm
LTE	long term evolution
MHz	megahertz
MIB	management information base
OFDM	orthogonal frequency division multiplexing
OSSI	operation(s) support system interface
PLC	1) physical layer link channel; 2) PHY link channel
PNM	proactive network maintenance
PSD	power spectral density
Q	quadrature (imaginary)
QAM	quadrature amplitude modulation
RBW	resolution bandwidth
RF	radio frequency
RxMER	receive modulation error ratio
SC-QAM	single carrier quadrature amplitude modulation
SCTE	Society of Cable Telecommunications Engineers
SNMP	Simple Network Management Protocol
TDR	time domain reflectometer
TFTP	Trivial File Transfer Protocol
TV	television
UHF	ultra high frequency

11. Bibliography & References

[Ref. 1]: Data-Over-Cable Service Interface Specifications DOCSIS 3.0 Operations Support System Interface Specification CM-SP-OSSIV3.0-I20-121113 (Cable Television Laboratories)

[Ref. 2]: Williams, Tom. "Correlating Return-Band Impulsive Noise Measurements from Houses with Sheath Current Induction Test Results," 1999 NCTA Technical Papers

[Ref. 3]: Data-Over-Cable Service Interface Specifications DOCSIS® 3.1 Cable Modem Operations Support System Interface Specification CM-SP-CM-OSSIV3.1-I16-190917 (Cable Television Laboratories)

[Ref. 4]: DOCSIS® Best Practices and Guidelines PNM Best Practices: HFC Networks (DOCSIS 3.0) CM-GL-PNMP-V03-160725 (Cable Television Laboratories)

[Ref. 5]: Data-Over-Cable Service Interface Specifications DOCSIS® 3.1 Physical Layer Specification CM-SP-PHYv3.1-I17-190917 (Cable Television Laboratories)

[Ref. 6]: Data-Over-Cable Service Interface Specifications DOCSIS® 3.1 CCAP™ Operations Support System Interface Specification CM-SP-CCAP-OSSIV3.1-I18-200610 (Cable Television Laboratories)

[Ref. 7]: Campos, Alberto; B. Currivan; C. Moore; T. Williams; “Upstream Cable Echoes Come In Two Flavors” *CED* magazine (on-line), February 2010. An enhanced version of the article was included in Appendix V of DOCSIS® Best Practices and Guidelines Proactive Network Maintenance Using Pre-equalization CM-GL-PNMP-V02-110623 (Cable Television Laboratories).

[Ref. 8]: PNM Best Practices Primer: HFC Networks (DOCSIS® 3.1) CM-GL-PNM-3.1-V01-200506 (Cable Television Laboratories)

12. Appendix

The following Python code is provided as a demonstration of how to retrieve and plot FBC data from a cable modem.

The script `getFbcData.py` should be called after a capture is initiated on a modem by setting the `docsIf3CmSpectrumAnalysisCtrlCmdEnable` object to true. The results are saved into the specified CSV file that can then be imported into a spreadsheet application such as Excel or be used by the `showFbcData.py` script.

12.1. File: `getFbcData.py`

Usage: `getFbcData.py` <SNMP community string> <cable modem IP address> <filename>

```
#!/usr/local/bin/python3
#
# THIS SOFTWARE IS PROVIDED BY THE AUTHORS AND CONTRIBUTORS "AS IS" AND
# ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
# WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
# DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE LIABLE FOR
# ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
# (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
# LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND
# ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
# (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
# SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
#
# Collects docsIf3CmSpectrumAnalysisMeasAmplitudeData from cable modem and save to
# CSV file
#
# This script requires the following python3 packages have been installed
#   pysnmp
#
# Author: James Medlock, jmedlock@akleza.com
#

import sys, os, getopt
import plotly.express as px
from pysnmp.hlapi import *

# Define FBC SNMP OIDs used
docsIf3CmSpectrumAnalysisMeasAmplitudeData=".1.3.6.1.4.1.4491.2.1.20.1.35.1.2"

# Convert 2's complement 16 hex value to decimal
def twos(val):
    x = int(val, 16)
    if x > 32767:
        x = x - 65535
    return x

def main():
    ampdata = []
    for (errorIndication, errorStatus, errorIndex,
        varBinds) in bulkCmd(SnmpEngine(),
        CommunityData(str(sys.argv[1])),
        UdpTransportTarget((str(sys.argv[2]), 161)),
        ContextData(),
        0, 2,
        ObjectType(ObjectIdentity(docsIf3CmSpectrumAnalysisMeasAmplitudeData)),
        lexicographicMode=False):

        if errorIndication:
            print(errorIndication)
            break
        elif errorStatus:
            print('%s at %s' % (errorStatus.prettyPrint(),
                errorIndex and varBinds[int(errorIndex)-1][0] or '?'))
            break
        else:
            for varBind in varBinds:
```

```

        ampdata.append(varBind[1].prettyPrint()[2:])

if len(ampdata) <= 1:
    print("No FBC data available")
    sys.exit()

# Decode each AmplitudeData row and write frequency,amplitude values to output file
with open(sys.argv[3], "w") as outF:
    for x in ampdata:
        centerfreq=int(x[0:8],16)
        freqspan=int(x[8:16],16)
        numbins=int(x[16:24],16)
        binspacing=int(x[24:32],16)
        resbw=int(x[32:40],16)

        startfreq = centerfreq - (numbins / 2 * binspacing)

        i = 0
        bin = 0
        while i < numbins*4:
            freq = startfreq + (bin * binspacing)
            j = 40 + i
            hexvalue = x[j:j+4]
            ampvalue = twos(hexvalue) / 100.0

            print(str(freq) + "," + str(ampvalue), file=outF)

            i += 4
            bin += 1

if __name__ == "__main__":
    if len(sys.argv) < 4:
        print("Usage: getFbcData.py <SNMP community string> <cable modem IP address> <filename>")
        sys.exit()

    main()

# End of file

```

12.2. File: showFbcData.py

Usage: showFbcData.py <filename>

```

#!/usr/local/bin/python3
#
# THIS SOFTWARE IS PROVIDED BY THE AUTHORS AND CONTRIBUTORS "AS IS" AND
# ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
# WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
# DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE LIABLE FOR
# ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
# (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
# LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND
# ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
# (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
# SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
#
# Display FBC data stored in CSV file
#
# This script requires the following python3 packages have been installed
#   plotly
#   pandas
#
# Author: James Medlock, jmedlock@akleza.com
#
# Usage: showFbcData.py <filename>
#

import sys
import csv
import plotly.express as px

def main():
    xdata = []

```

```

ydata = []
with open(sys.argv[1]) as csvfile:
    csvData = csv.reader(csvfile, delimiter=',')
    for row in csvData:
        xdata.append(float(row[0]) / 1000000)
        ydata.append(float(row[1]))

fig = px.line(x=xdata, y=ydata,
              labels={'x': 'Frequency (MHz)', 'y': 'Amplitude (dB)'},
              title='DOCSIS Cable Modem Full Band Capture')
fig.update_layout({'plot_bgcolor': 'rgba(0, 0, 0, 0)', 'paper_bgcolor': 'rgba(0, 0, 0, 0)'},
                  title={'x': 0.5, 'xanchor': 'center'})
fig.update_traces(hovertemplate='Frequency: <b>{x:.2f} MHz</b><br>Amplitude: <b>{y:.2f} dB</b>',
                  line_color='rgb(68,114,196)')
fig.update_xaxes(showline=True, linewidth=2, linecolor='black', gridcolor='rgb(209,209,209)')
fig.update_yaxes(showline=True, linewidth=2, linecolor='black', gridcolor='rgb(209,209,209)')
fig.show()

if __name__ == "__main__":
    if len(sys.argv) < 2:
        print("Usage: showFbcData.py <filename>")
        sys.exit()

    main()

# End of file

```


Full Scale Deployment of PMA

Lessons Learned from Deploying the Profile Management Application System at Scale and Considerations for Expanding the System Beyond OFDM

A Technical Paper prepared for SCTE•ISBE by

Maher Harb

Director, Data Science
Comcast
1800 Arch Street, Philadelphia, PA 19103
267.260.1846
maher_harb@comcast.com

Bryan Santangelo

Executive Director, Data Eng and Science
Comcast
1800 Arch Street, Philadelphia, PA 19103
918.640.8936
bryan_santangelo@comcast.com

Dan Rice

Vice President, HFC Architecture
Comcast
1401 Wynkoop St Ste 300, Denver, CO 80202
720.512.3730
daniel_rice4@comcast.com

Jude Ferreira

Principal Engineer
Comcast
1800 Arch Street, Philadelphia, PA 19103
215.286.4070
jude_ferreira@comcast.com

Table of Contents

Title	Page Number
Introduction	3
1. Overview of System and Algorithms	5
2. Network Visibility	7
2.1. Home level view	8
2.2. Interface level view	10
2.3. Notifications.....	11
3. Next Generation Algorithms	11
4. Upstream DOCSIS 3.0 Profile Management	13
4.1. Overview	13
4.2. How D3.0 US PMA works.....	16
4.3. Why D3.0 US Optimization increase network capacity?	18
4.4. The Reinforcement Learning approach to US PMA.....	20
5. Conclusion	21
Abbreviations.....	22
Bibliography & References	23

List of Figures

Title	Page Number
Figure 1 - COVID network impact on traffic peak times.	4
Figure 2 - COVID network impact on bandwidth demand.	4
Figure 3. The PMA System Architecture.	5
Figure 4. Example of PMA profiles constructed for a set of 20 devices.	7
Figure 5. Screenshot of the DS PMA home view dashboard.	8
Figure 6. Interface level view dashboard for DS PMA.....	10
Figure 7 - Key results from the experiment in which modulation mapping thresholds were relaxed.	13
Figure 8 - Schematic demonstrating how upstream Profile Management would enable additional capacity in legacy low split HFC.	14
Figure 9. Peak COVID High Utilization Example.	15
Figure 10 - US Dashboard example for a HUB of 5 new CMTSS managed by the US D3.0 PMA solution.	16
Figure 11 - The profile configuration template for 6.4 MHz-wide channel.	17
Figure 12 - D3.0 PMA System Comes Together.....	18
Figure 13 - Anatomy of a FEC codeword showing example of different codeword configurations for the D3.0 US long data grant.....	19
Figure 14 - Schematic of the Reinforcement Learning system for US PMA.....	20

List of Tables

Title	Page Number
Table 1 - Minimum MER values that support the corresponding modulation from DOCSIS 3.1 specification.	12

Introduction

In 2019, Comcast developed a Profile Management Application (PMA) system for generating and transacting D3.1 downstream (DS) profiles tailored to the conditions of each Orthogonal Frequency Division Multiplexed (OFDM) channel in its network. The approach, machine learning algorithms and system architecture were described in a previous SCTE technical paper [1]. The initial plan for this follow-up paper was to focus on Comcast's PMA deployment journey, the success of which is evidenced by thousands of Cable Modem Termination Systems (CMTSs) managed by the PMA, yielding greater than 20 Tbps of added downstream (DS) capacity to the network.

With the onset of the COVID-19 crisis, some of that focus shifted, in lockstep with the shift of the U.S. and worldwide workforce from office to home. Figure 1 shows the 32% increase in upstream (US) traffic, post-COVID, and the shift in peak times for DS traffic from 9:00 PM to 7:30 PM, and from 9:00 PM to 8:00 AM & 6:00 PM for US traffic. Figure 2 shows the bandwidth demand growth, around time of the COVID crisis (Spring of 2020), for US traffic (black curve) and DS traffic (sky blue curve). With work-from-home traffic increasing on the network, and because of the earlier implementation of the PMA, the DS capacity was available, and the network was able to easily scale to the significantly increased demand.

The US is a different story. As a fraction of the total available spectrum, and even as it is being industrially widened from sub-split to high-split configurations, the fact remains that US capacity is a more difficult challenge. Commencing with shelter-at-home requirements, US traffic grew sharply, seemingly overnight. Comcast has publicly shared data on the increases in traffic scale since COVID started [2-4], along with transparency about the level of investment and technological attention that prepared us for “Black Swan” scenarios like a pandemic. This enabled more effective management of the additional traffic growth delivered over the Data Over Cable Service Interface Specification (DOCSIS) broadband network [5]. As this paper will ultimately show, by adding an upstream PMA focus to the existing PMA suite, we were able to boost upstream capacity by 36%, from 86 Mbps to 117 Mbps.

Given these extraordinary circumstances, with the COVID crisis in full swing, and with the shift in internet usage, we refocused this paper to share our accelerated efforts in developing and deploying PMA for the US using DOCSIS 3.0 technology. Fortunately, the technology that was brought to bear on the challenges of US capacity was already under development. The effort was shaped by the early concepts found in CableLabs member publications from the late '90s [6] into the early 2000s. Combined with state-of-the-art methods, including scaled cloud-based compute, and machine learning techniques such as Reinforcement Learning (RL), we were able to ensure system stability and optimize the network bandwidth as spectral conditions and demand changed.

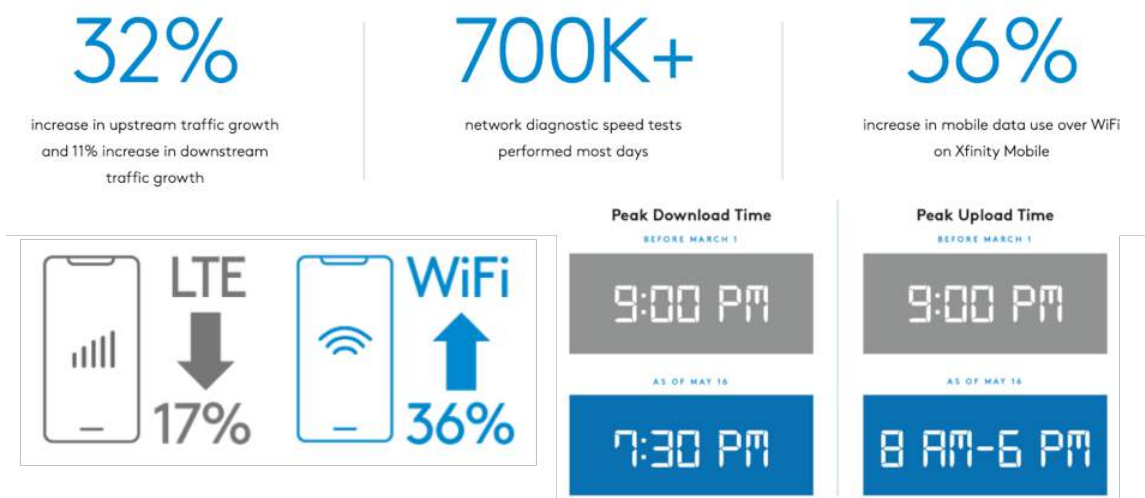


Figure 1 - COVID network impact on traffic peak times.




Figure 2 - COVID network impact on bandwidth demand.

Note that at the start of 2020, the US D3.0 profile management was under development as part of a holistic capacity management and long-range access network architecture plan, along with other efforts, such as passive hybrid fiber-coax (HFC) with deeper fiber, virtualized CMTS with remote PHY nodes, and spectrum augmentations, such as high-split and FDX. The D3.0 US PMA effort was initially targeted at optimizing the upstream spectrum, in conjunction with the deployment of additional US channels. The additional US channels were located in parts of the spectrum that would be difficult to manage without an autonomous modulation profile optimization system, given known and persistent levels of ingress. The D3.0 US profile management was also intended for those network segments without fiber deep and spectrum upgrades, to help offset the timing risks related to developing and deploying new technology by deferring or eliminating investment in legacy technologies.

The paper is organized as follows: Section 1 recaps our 2019 Expo paper by providing an overview the PMA system and its algorithms and describing the current state of the DS algorithm. Section 2 presents analyses and dashboards created to track system performance & health. Section 3 presents our vision for

how the PMA system is expected to evolve in the future. Section 4 is dedicated to describing the US PMA system. Section 5 concludes with the lessons learned from the PMA deployment journey.

1. Overview of System and Algorithms

The PMA system architecture and DS algorithms have been described at length in the previous SCTE technical paper [1]. A brief recap and description of the prior work is included here for professional context. The PMA system is composed of four separate components, shown in Figure 3: Data Collector, Data Storage, Analytics Engine, and Configuration Manager.

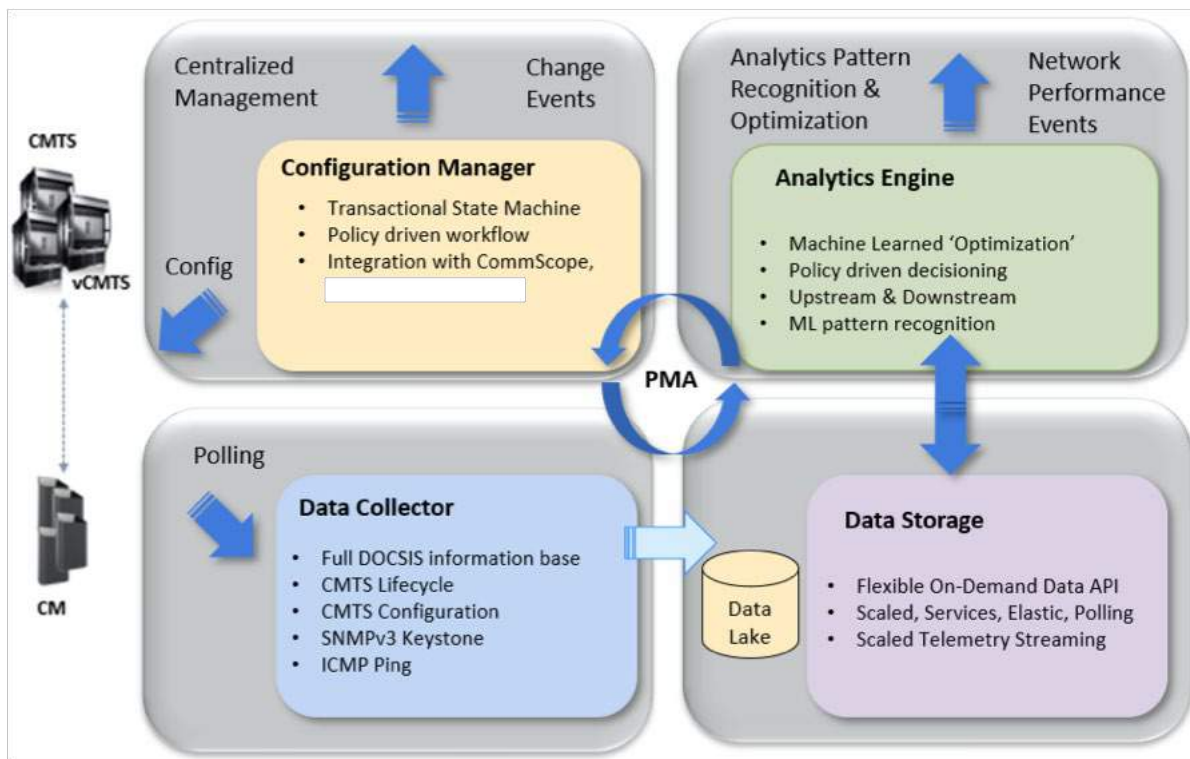


Figure 3. The PMA System Architecture.

The Data Collector is responsible for collecting telemetry data from CMTSs and gateway devices. The data is polled at different frequencies that range from every 5 min to hourly, and was designed to constitute a “comprehensive poller,” enabling applications beyond the scope of PMA. From a PMA perspective, the data needed to support the construction of OFDM profiles falls into the following categories:

- **Network topology:** Establishes linkage between device, OFDM channel, and CMTS.
- **Configuration model:** Provides characteristics of the OFDM channel, e.g. number of subcarriers, subcarrier width, frequency range, position of exclusion bands, etc.
- **CMTS type:** Provides make, model, hardware & software versions of a given CMTS.

- **Telemetry:** Retrieves Modulation Error Ratio (MER), Forward Error Correction (FEC), signal, and traffic measurements from devices, and channel utilization measurements from CMTSs. This category constitutes the largest bulk of the data, given that MER spectra are measured at a per-device OFDM subcarrier resolution, with 4096 data points for each MER sample for each device.

The Data Storage is primarily comprised of a public cloud-based data lake, where the polled data listed above land in raw (unprocessed) form.

The Analytics Engine (AE) is a machine learning pipeline that uses the data to construct OFDM profiles suitable for use by the devices in the network—given spectral conditions measured over certain time windows. At its core, constructing profiles is a type of optimization problem in which the stated objective is to maximize channel capacity and minimize codeword error rates, subject to certain constraints. Thus, the problem contains an inherent trade-off between improving robustness and increasing network capacity, since reducing error rates is achieved by opting for lower modulation levels, at the expense of reduced channel capacity.

The constraints are dictated by the CMTS hardware and software versions, as different CMTSs support different numbers of profiles per OFDM channel. Within the construct of a profile, they may also support different numbers of modulation exception zones (segments), as well as imposing additional constraints on the attributes of a segment (e.g. segment width).

Algorithmically, the AE uses hierarchical clustering—a type of unsupervised machine learning algorithm—to group together devices that share common noise characteristics and assign them a common modulation profile. Additional smoothing algorithms are applied post-clustering, to reshape the segments according to given constraints. In the current version of the algorithm, the clustering objective function is designed to maximize capacity around a statistical decision boundary.

FEC rates are considered, indirectly, by imposing additional constraints on the mapping from MER values to modulation levels (e.g. a MER value > 27 dB supports 256-QAM at maximum). As an example, the plot in Figure 4 shows MER measurements alongside the constructed profiles on a dual y-axis plot. Since spectral conditions vary over time, multiple MER samples are captured over a time window dictated by AE policy. For each panel (device) we show 3 curves characteristic of the variation in MER: the max level (dark gray curve), the min level (light gray curve), and the 10th percentile (red curve). Also, per policy, it is the 10th percentile that is fed to the algorithm as conservatively representative of the device's MER state. The constructed profiles are overlaid on the plots and follow the scale of the right y-axis. In this specific example, the CMTS allows 4 profiles per OFDM channel, 4 segments per profile, and a segment width that is a multiple of 1 MHz. Profiles 1-3 are overlaid in yellow, blue, and green colors, respectively on the devices that are assigned to each of the 3 profiles. Profile 0 (not shown) is the control profile and is set to a flat 256-QAM by AE policy. Note that the impairments shown are generated in the lab and applied to select devices. Because of the CMTS-imposed limitation of 4 exception zones (segments), the algorithm overcompensates for the V-shaped impairment exhibited in the MER spectra of device #5.

Lastly, the Configuration Manager (CM) is responsible for transacting profiles generated by the AE. The output from the AE defines profiles according to a standardized intermediate JSON format that is agnostic to the CMTS make and model. The CM converts the output to commands that are specific to the CMTS. The CM is also responsible for validating the profiles, deciding on whether to reject or accept the AE recommendations, scheduling the transacting of the profiles according to a policy that defines allowed maintenance dates/times, and performing pre- to post-transaction checks to confirm that the configuration was successfully applied.

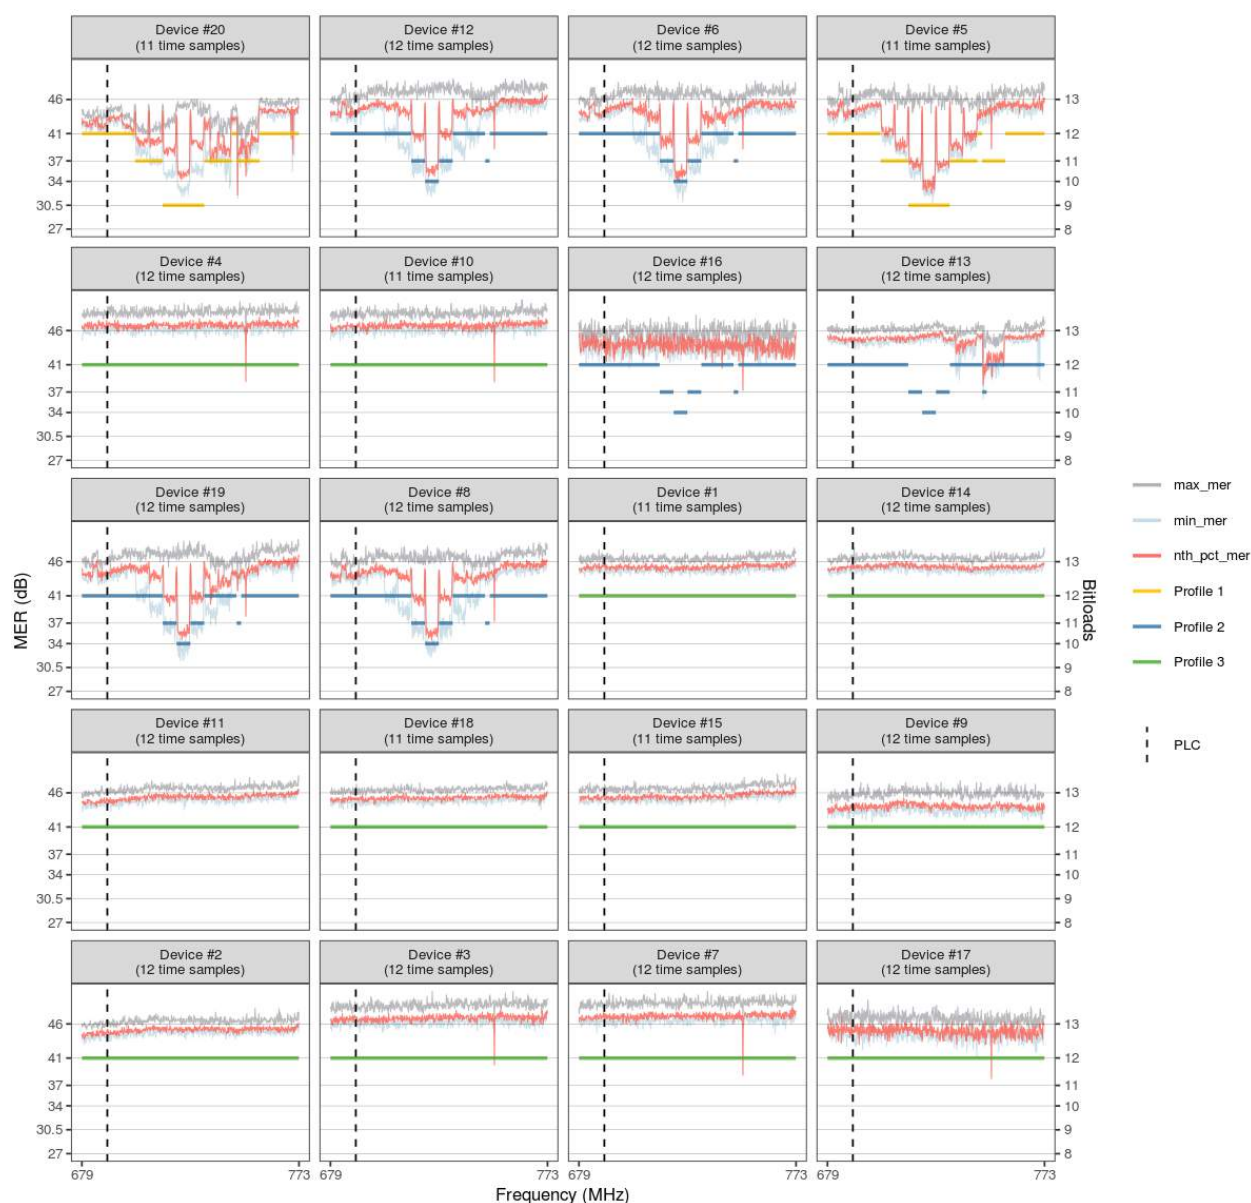


Figure 4. Example of PMA profiles constructed for a set of 20 devices.

2. Network Visibility

A prerequisite to deploying PMA was achieving a high level of visibility into the network, in order to respond to issues that may arise as deployment ramps up. Thus, this requirement opened a path to an independent effort on dashboarding and automated notifications. While the PMA's Data Collector is exhaustive in the breadth of data it collects, the focus of this effort was to provide insight and visibility into metrics that are PMA-related or perceived to be PMA-related. We start by introducing the basic design criteria for the PMA dashboards:

- The overall health of the PMA system is to be monitored through a common “home view” that constitutes an entry point to all other dashboards.

- One must be able to drill down from the “home view” to an “OFDM channel” detailed view.
- All dashboards must display near-real time information (limited by the Data Collector’s polling frequencies).
- Dashboards must be informative in the operational sense, i.e. highlighting issues and allowing operations to take appropriate actions to remediate.

2.1. Home level view



Figure 5. Screenshot of the DS PMA home view dashboard.

Given the design criteria, we next introduce some of the key metrics displayed on the home view dashboard shown in Figure 5. The home view aggregates all metrics by division—the highest-level organizational structure within Comcast’s network. This view is meant as a first entry point to monitor the health of the PMA system. The “Lowest 100 Ranked Interfaces” table brings to attention interfaces (OFDM channels) that are low performing, from a PMA perspective and, thus require further investigation. Clicking on an interface name opens the interface-level dashboard with additional details.

Note that names of divisions and interfaces shown in the dashboard screen capture were anonymized. The metrics shown on the home view are:

- **Point-in-time Metrics:** Latest metrics on health and performance of the system that include:
 - **Capacity gain** (34.3% for Division A in Figure 5): The capacity gain attributable to PMA as measured relative to a 256-QAM baseline. Note that the capacity metric represents an instantaneous value calculated from the actual traffic distribution across different OFDM profiles as captured from telemetry in 100s of thousands of service groups across the network.
 - **Raw gain** (6020 Gbps for Division A in Figure 5): Same metric as above but calculated in units of absolute Gbps.
 - **Number of CMTSs** (587 for Division A in Figure 5): Number of CMTSs managed by the PMA system at the time of this snapshot.
 - **Traffic on Profile 0** (2.1% for Division A in Figure 5): Since Profile 0 is configured as 256-QAM or 64-QAM by policy, ideally no devices would use Profile 0. Thus, this metric is indicative of a combination of the health of the system and the freshness of the configured profiles. If spectral conditions degrade and data profiles are not updated in a timely manner, more traffic will flow on Profile 0.
 - **CM success rate** (91.0% for Division A in Figure 5): Percent of CMTSs that were successfully configured with updated profiles during the last configuration window, which is set by policy.
 - **Excluded modems** (0.8% for Division A in Figure 5): Percent of devices with severe impairments that were excluded from the clustering algorithm. The rationale for device exclusion is to avoid wasting a customized profile on severely impaired devices that require field work to effectively mitigate. The current criteria for excluding a device requires that the mean modulation that a device supports is less than the profile 0 capacity (256-QAM or 64-QAM).
- **Lowest 100 ranked interfaces:** The aggregate measures described above may conceal severe issues affecting only a small number of OFDM channels. Hence, this table ranks the 100 channels that require most attention. The ranking is based on a combination of metrics that include traffic on Profile 0, number of devices experiencing OFDM partial service issues, and total number of devices on the channel. In addition, this table provides an entry to detailed interface-level dashboard by clicking on the channel name.
- **Traffic distribution:** This is a time series showing the distribution of traffic across profiles aggregated for all OFDM channels in the topology context. The capture shows that >90% of traffic flows through Profile 3—the highest capacity profile. It also reveals some cyclical behavior with increased traffic on Profile 2 during peak usage times. Note that PMA configures profiles based on device clustering according to the MER characteristics. Even though the PMA assumes that each device should be assigned a specific profile, the ultimate assignment of profiles is left to the CMTS, according to its internal profile selection function as a policy. The fact that most of the traffic flows on Profile 3 indicates that the mapping between MER values and modulation levels is conservative. Section 3 discusses capturing additional capacity gain by

adjusting the AE policy thresholds.

- **Age of profiles:** A distribution of profile age shows that most interfaces were updated within the last day (column labeled 0 days). This view allows monitoring profiles that may become stale. The current update frequency is daily. However, the AE will recommend profile updates only if there was a change in the spectral conditions that warrants updating the existing profile. Hence, it is expected that the age distribution will show profiles that are older than 0 days.

2.2. Interface level view

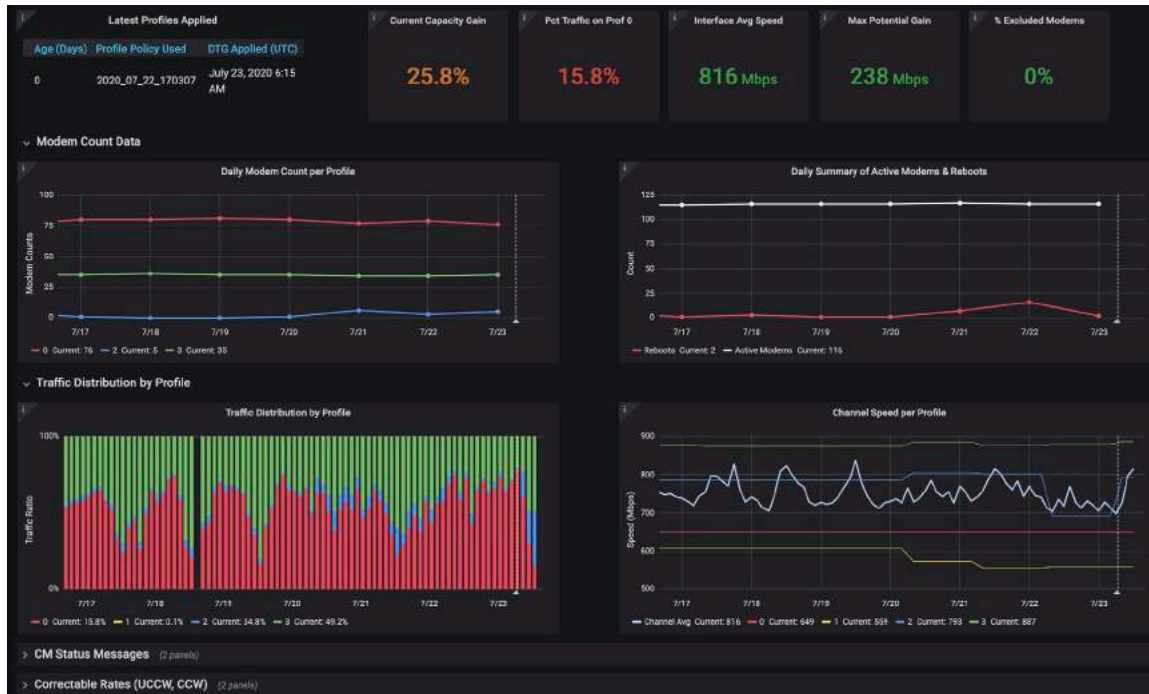


Figure 6. Interface level view dashboard for DS PMA.

The interface (OFDM channel) level dashboard offers a highly detailed view of the network dynamics occurring on that interface. A screenshot of the interface dashboard is shown in Figure 6. It shows point-in-time metrics as well as time series metrics covering various dimensions. This view shows device counts and reboots, traffic distribution, and channel speed as function of time. Other views including the CM status message statistics and error rates are “collapsed in” (not shown in screen capture) for conciseness. The view contains several sections covering the following areas:

- **Point-in-time metrics:** These are similar to the home-level view of point-in-time metrics but aggregated at the interface-level.
- **Modem counts data:** Counts of devices by profile, total device count and count of device reboots.
- **Traffic distribution:** Traffic distribution across profiles and the speed of each profile.

- **Modem CM-Status Messages:** Breakdown of D3.1 CM-status messages reported by devices related to the health of the OFDM channel.
- **FEC error rates:** Time history of the correctable and uncorrected codeword error rates.
- **Traffic volume data:** Time history of traffic volume (unicast octets) as reported both by CMTS and devices.

Note that similar dashboards were developed for the US PMA. Select views of these are shown in Section 4.

2.3. Notifications

Because networks tend to degrade over time from an impairment perspective, if not maintained, we developed methods to reduce the impact of the network on the customer experience. In addition to the operational views described above, impairment detection algorithms are provided as a core part of the architecture. Notifications from these are delivered across a messaging bus to other Comcast OSS tools, to ensure that technicians are dispatched to the right hubs, network segments and homes to remediate issues. The following data are provided as part of the notifications, to assist with event prioritization and triangulation:

- Interface details such as Physical Link Channel (PLC) location, start/end frequencies, total number of impacted cable modem counts.
- List of impacted cable modems with severity of impairments.
- Interface impairment ranking at the national, divisional, and regional levels.
- Reference to the API for historical data stored in the data lake related to the event.
- Reference for the API to enable fix agents to collect real-time on demand data, to confirm that the issue is still present, and to assist in isolation and to confirm mitigation.
- Other data sources to enrich the event, such as the mobile wireless carriers that overlap with the OFDM channel.

Notifications for severely impacted interfaces are also sent via email to divisional engineering and operations teams. Additional details, such as device/interface RxMER images, are posted to internal collaboration platforms. Note that another Comcast-authored SCTE paper describes the pattern detection associated with these notifications [7].

3. Next Generation Algorithms

The current PMA algorithm predicts a ~30% increase in network capacity, relative to a flat 256-QAM baseline, on average. Figure 5 shows that the capacity gain as measured from real-time traffic is around ~35%. This gap between predicted gain and realized gain is due to the internal profile selection function of the CMTS, based on CM-STATUS messages. The internal profile selection effect is clearly visible in the traffic distribution by profile shown in Figure 5, in which more than 90% of the traffic is observed to flow through Profile 3—the highest capacity profile. The logic for internal profile selection function varies across vendors, but the general idea is the same: the CMTS regularly sends test codewords to devices on all configured profiles and determines, based on the encountered errors, the highest capacity profile a device is capable of using. The CM responds with CM-STATUS messages descriptive of that specific device's perspective on the performance. These messages are further processed through proprietary selection algorithms with controls for the operator to optimize. The fact that most devices are able to use a higher capacity profile than recommended by the PMA algorithm indicates that the

thresholds adopted for converting MER to modulation (shown in Table 1) are somewhat conservative. This added resiliency against errors is a known feature of the low density parity check (LDPC) error correction algorithm. While quantifying the LDPC benefit is not straightforward, prior studies suggest that it provides an additional 3-6 dB improvement over Reed Solomon—its D3.0 counterpart [8] Next, we discuss how to take advantage of the superior performance of the LDPC to capture additional capacity gain.

Table 1 - Minimum MER values that support the corresponding modulation from DOCSIS 3.1 specification.

MER Threshold (dB)	Modulation efficiency
0	0
9	2
15	4
21	6
24	7
27	8
30.5	9
34	10
37	11
41	12

The PMA pipeline offers several algorithmic tuning knobs that are configured by policy. One such knob is a global offset in dB applied to the thresholds listed in Table 1. We experimented with the knob by relaxing the threshold for 2 production CMTSS, while tracking key performance metrics against a control CMTS group within the same site. The results of the experiment shown in Figure 7 reveal that an additional 10% increase in capacity is garnered by relaxing the thresholds by 3 dB, without introducing a negative impact on the codeword error rates. In all panels, the control group is shown in red and the experiment group in light blue. The first 3 vertical dashed lines correspond to the points in time at which the thresholds were adjusted by 1, 2, and 3 dB respectively. For the 4th vertical dashed line, the adjustment was maintained at 3 dB (i.e. no change over previous state). The top left panel in Figure 7 shows the capacity of the experiment group increased per adjustment, with a total increase of ~10% at 3 dB. The top right panel shows the traffic on Profile 0 (256-QAM) shows similar levels across the 2 groups, indicating that shifting the thresholds is not causing data profiles to become unusable by devices. The bottom left panel shows the correctable error rates increased for the experiment group following the first adjustment by 1 dB. The bottom right panel shows the uncorrectable error rates reveal no trend related to the adjustments (notice that the uptick towards the end of timeline was experienced by both experiment and control groups).



Figure 7 - Key results from the experiment in which modulation mapping thresholds were relaxed.

While the results shown in Figure 7 are encouraging and point to the possibility of capturing additional gains by tweaking the global policy, we opted not to adopt such strategy for optimizing capacity gains. The main motivation is the fact that plant conditions vary across the network. A one-threshold-fits-all policy may mask problems in areas where impairments exist and where certain population of devices may benefit from some added robustness (even moving the threshold in the opposite direction, to sacrifice capacity for robustness). Hence, our current effort focuses on developing a Machine Learning (ML) control system that offers much greater flexibility in configuring the profiles, including adjusting the modulation thresholds down to within a specific profile exception zone (segment). Other policy attributes may similarly be modified, such as the statistical aggregations used across time and frequency, before and after clustering of modems.

One methodology that shows promise in this realm is reinforcement learning (RL). In RL, the ML agent “learns” an optimal policy by interacting with the environment. The outcome is akin to allowing the agent to dynamically modify the MER mapping thresholds, or other policy attributes, per OFDM channel-profile-exception zone (segment) and based on feedback in the form of the FEC error rates encountered by devices. We are currently in the midst of building a RL solution for US PMA. As will be shown in Section 4, US PMA has a limited action space, compared to DS PMA, and therefore it offers an opportunity to experiment with and refine the solution with the expectation that these methods will be subsequently adapted to be used for DS PMA and D3.1 US OFDMA PMA.

4. Upstream DOCSIS 3.0 Profile Management

4.1. Overview

The goal for the D3.0 US PMA solution was to add enough US capacity to allow at least 1 year of bandwidth demand growth, based on forecasted Compound Annual Growth Rates (CAGR), without having to segment nodes using legacy technology. Our engineering models calculate capacity based on MAC layer data rates, net of all physical layer and time-based overhead. Based on these models and progress to-date, using both D3.0 upstream PMA and adding the 5th & 6th channels, we have increased the

upstream MAC layer data rate capacity from approximately 86 Mbps to approximately 117 Mbps, an increase of 36%, exceeding a year's worth of upstream growth even if COVID traffic levels remain as described in Figures 1 & 2.




Figure 8 - Schematic demonstrating how upstream Profile Management would enable additional capacity in legacy low split HFC.

The profile management optimization process uses all the different knobs available in the upstream channel and modulation profile configuration to increase capacity, where channel quality allows it (see Figure 8). The top panel shows the addition of the 5th and 6th upstream channels; together, these are responsible for increasing capacity by approximately 15%. We are migrating to a standard channel plan across our sub-split nodes that enables a channel width of 3.2 MHz for the 5th US channel and 1.6 MHz for the 6th US, where both represent SC-QAM channels. Figure 8's bottom panel is a schematic highlighting the tradeoff between increasing capacity and increasing robustness. US PMA configures the knobs shown in the schematic to achieve the right balance. These tuning knobs are responsible for additional 20% capacity across all channels, above the robust defaults in use across the network in Q1 2020. Also, when needed, the solution can increase Forward Error Correction (FEC) and use more robust modulation to mitigate potential customer experience issues attributable to noise and ingress. These configuration updates are all done autonomously on the production network today. However, when the COVID-induced increase in upstream traffic began, the US automated execution flows had not been completed. Specifically, the configuration manager (CM) features were still under development.

As the capacity impact became apparent, the development team came together to work closely with the field operations and engineering teams, to calculate and statically configure more efficient modulation profiles with AE. The intent was to provide capacity relief and prove out the benefits of optimization algorithms on the production network, while simultaneously accelerating the development of the closed-

loop autonomous system. The initial static deployments based on the profiles calculated by the AE, as the shelter-at-home orders peaked – more so in markets saturated with high-tech corporations -- had a significant positive effect on capacity. One example of an extreme upstream congestion case is shown in Figure 9. The top plot shows octet utilization for 4 channels, and the bottom plot the minislot and octet utilization of one of the optimized channels. Profile upgrades were done on the 1st and 2nd channels but not the 3rd and 4th channels. Consequently, approximately 10% more peak data is seen going through network with the initial efficient profiles. In this example, the minislot utilization and the octet utilization were at the maximum level for extended periods of time. When the static, more efficient profiles were applied to the 2 higher spectrum channels, the capacity was increased by 17% relative to the capacity of the two lower spectrum channels, which were not modified until the fully autonomous CM was available. Additionally, for the same level of minislot utilization, 10% more Mbps (octet utilization) were able to be sent by customers through the network than were previously transmittable. Correspondingly, the minutes of time at maximum utilization per day were subsequently reduced. Adding additional channels to further augment capacity on top of this example, along with other techniques, enabled a very quick response to and resolution of capacity hotspots.




Figure 9. Peak COVID High Utilization Example.

Figure 10 shows an example of one of the operational dashboard's tracking system performance for the upstream PMA solution. Notice the activity in the profile distribution plots (identifiable by the colored bars around 7/18). This activity corresponds to the early stages of introducing a new CMTS, because the control system requires a few update cycles before converging to steady-state, at which time which most channels operate at the highest capacity profile (profile index 251). For reference, as additional CMTSs are added to the system, they are initialized to start on static, very robust profiles. These robust profiles are intended to combat a common transient noise source on the network in the lower 2 spectrum channels.

Additionally, and as noted earlier, a very robust but efficient modulation profile had also been statically deployed on the upper two spectrum channels, as described in the 2018 SCTE paper [9]. After accelerating the CM development and moving through trial and deployment, the automated system was turned on. It iterated through a set of modulation profiles, in small, low risk increments, and based on the channel's quality metrics. The metrics include correctable and uncorrectable codeword errors, signal-to-noise-ratio (SNR), partial service statistics, and minislot utilization. Over a period of approximately 24 hours, the modulation profiles were modified until they reached a steady-state based on the set control system strategy. The result is approximately 98% of the upstream channels are running effectively, without codeword errors, and around 2% are running on a profile addressing US noise sources such as transient switching power. The 21.2% capacity gain shown on the dashboard for these CMTSs is consistent across our network, approximately doubling the capacity gains of the example shown in Figure 9. The upshot is that the system autonomously optimizes all upstream channels, not just those in the higher quality spectrum.




Figure 10 - US Dashboard example for a HUB of 5 new CMTSs managed by the US D3.0 PMA solution.

In addition to other key performance indicators, such as system settling time, capacity gain and error rates, there is the profile stability rate. Also shown on the dashboard is a distribution of how often an upstream channel has a profile update, highlighting remarkable network stability. In total, 98% of US channels require a profile update less than 1 time per week; 99% of US channels require a profile update less than 3 times a week; and < 0.5% of channels require profile updates more than once per day.

This stability also highlights the promise of D3.1 OFDMA profile management as the modem population grows, or the upstream spectrum available allows us to achieve the benefits of OFDMA. The authors plan to provide a future update on the OFDMA algorithm under development in 2020. What's exciting so far is that network quality is superior to that required for the best D3.0 capacities, and the profile management system will be able to take advantage of the higher order modulations available with D3.1.

4.2. How D3.0 US PMA works

The PMA system, as described previously, was extended to implement US D3.0 PMA functionality. The capacity is increased (or decreased) in small steps, while errors are fixed proportionally or predictively, either by increasing robustness or detecting transient noise indicators. When network issues are detected, operations staff is notified of network issues to drive remediation, ensuring maximum capacity is sustained and yielding operational savings. This is a “business as usual” step that happens regardless of whether the customer experience is being degraded due to packet loss.

The modulation profile capacities shown in the Figure 11 are based on compatible upstream channel configurations and channel widths. Similar templates exist for the narrower 3.2 MHz and 1.6 MHz channels. Figure 11 summarizes a subset of modulation profile attributes that must be set compatibly with the US channel attributes and other aspects, such as codeword size, preamble length, guard time, and

interleaver settings. For example: profile 251 uses a 97 bytes payload and a 2 bytes parity for the short data grant, a 247 bytes payload and a 4 bytes parity for the long data grant, and designates the cutoff between short and long data grants to be 5 minislots in length (the burst size). The station maintenance and unsolicited grant service interval usage code (UGS IUCs) are similarly optimized to achieve the efficient use of minislots and required robustness. Each template consists of 25 profiles, constructed to comprehensively sample the parameter space, along the modulation and FEC regime dimensions. For example, profile 251 exhibits the highest modulation (256-QAM) and least robust profile (meaning the profile with the lowest FEC overhead.)

	251	256	261	266	271
	25.6 Mbps QAM64 short: 97/2, burst=5 long: 247/4 SNR for 1% error rate = 22.8 dB	24.5 Mbps QAM64 short: 91/5, burst=5 long: 239/8 SNR for 1% error rate = 22.1 dB	23.3 Mbps QAM64 short: 105/10, burst=6 long: 229/13 SNR for 1% error rate = 21.2 dB	22.5 Mbps QAM64 short: 99/13, burst=6 long: 223/16 SNR for 1% error rate = 20.5 dB	20.7 Mbps QAM64 short: 99/13, burst=6 long: 121/16 SNR for 1% error rate = 20.3 dB
	252	257	262	267	272
	21.4 Mbps QAM32 short: 98/2, burst=6 long: 247/4 SNR for 1% error rate = 20 dB	20.5 Mbps QAM32 short: 92/5, burst=6 long: 239/8 SNR for 1% error rate = 18.8 dB	19.4 Mbps QAM32 short: 102/10, burst=7 long: 229/13 SNR for 1% error rate = 18.2 dB	18.8 Mbps QAM32 short: 96/13, burst=7 long: 223/16 SNR for 1% error rate = 17.5 dB	17.8 Mbps QAM32 short: 96/13, burst=7 long: 138/16 SNR for 1% error rate = 17.3 dB
	253	258	263	268	273
	17.2 Mbps QAM16 short: 91/2, burst=7 long: 247/4 SNR for 1% error rate = 16.6 dB	16.5 Mbps QAM16 short: 101/5, burst=8 long: 239/8 SNR for 1% error rate = 15.9 dB	15.6 Mbps QAM16 short: 91/10, burst=8 long: 229/13 SNR for 1% error rate = 15.2 dB	15.1 Mbps QAM16 short: 101/13, burst=9 long: 223/16 SNR for 1% error rate = 14.6 dB	14.3 Mbps QAM16 short: 101/13, burst=9 long: 138/16 SNR for 1% error rate = 14.4 dB
	254	259	264	269	274
	13 Mbps QAM8 short: 100/2, burst=10 long: 247/4 SNR for 1% error rate = 14.7 dB	12.4 Mbps QAM8 short: 94/5, burst=10 long: 239/8 SNR for 1% error rate = 14 dB	11.7 Mbps QAM8 short: 96/10, burst=11 long: 229/13 SNR for 1% error rate = 13.2 dB	11.3 Mbps QAM8 short: 90/13, burst=11 long: 223/16 SNR for 1% error rate = 12.5 dB	10.8 Mbps QAM8 short: 90/13, burst=11 long: 146/16 SNR for 1% error rate = 12.4 dB
	255	260	265	270	275
	8.7 Mbps QPSK short: 93/2, burst=14 long: 247/4 SNR for 1% error rate = 10.3 dB	8.3 Mbps QPSK short: 95/5, burst=15 long: 239/8 SNR for 1% error rate = 9.3 dB	7.8 Mbps QPSK short: 85/10, burst=15 long: 229/13 SNR for 1% error rate = 8.6 dB	7.6 Mbps QPSK short: 87/13, burst=16 long: 223/16 SNR for 1% error rate = 7.9 dB	6.7 Mbps QPSK short: 89/15, burst=17 long: 104/16 SNR for 1% error rate = 7.6 dB
(lower) Modulation <- -> (higher)	(efficient) <- -> FEC Regime --> (robust)				

Figure 11 - The profile configuration template for 6.4 MHz-wide channel.

Figure 12 describes how the whole PMA system came together. In the top schematic, theoretical models for noise are shown, as well as measured MER and traffic data that informs the construction of the US profiles. The bottom left schematic shows the result, which is a profile matrix configuration that spans a wide range of operating conditions. On the bottom right, system response is shown, in terms of error rates along with a defined policy that informs the selection of a suitable profile. This iterative process continues indefinitely as network conditions change. Note that the current profile update cycle is once every 6 hours. This value is configurable through policy and we are currently conducting A/B testing on a small subset of the CMTSs to discern if a higher update frequency (~hourly) further improves performance. In more detail:

- Models were developed for different transient and white noise channel models for each of the profiles.
- The range of network metrics were analyzed across the different spectrum locations, coming from over 50M+ DOCSIS devices, to understand the range of performance expected on the production network.
- Distribution of traffic packet and concatenated burst data was obtained to fine tune the codeword efficiency calculations.
- This data was fed into models of profile design, resulting in a set of static modulation profiles that cover a wide range of operating conditions and capacity yields.
- These profiles were tested in an automated lab system that automatically configures the profile on the upstream channel and injects the additive white Gaussian noise AWGN or transient noise, at increasing intensity, in increments of a fraction of a dB. Simultaneously, the system measures the CMTS upstream metrics for SNR and metrics related to codewords and traffic generators. The lab results were then compared to the theoretical models and matched to within +/- 1 dB. These lab and modeling results were then used to set the profile management application policy.
- The analytics engine, based on these models, analyzes the network data in real time. Based on statistics across time samples, and driven by policy, it selects the correct modulation profile for the current conditions. A policy, for example, could limit the operation to a subset of the matrix of profiles, such as limiting the low spectrum 6th channel to only operate in lower modulation transient noise profiles, or constraining the higher spectrum channels to only operate in the top 2 rows of the matrix, depending on channel minislots utilization levels.
- The analytics engine then closes the loop by continuing to measure network performance. It adjusts the modulation profiles as required, and based on network performance statistics intended to keep the errors at a healthy level, also based on policy.
- This closed loop control system is modifiable using Reinforcement Learning (see Section 4.4).

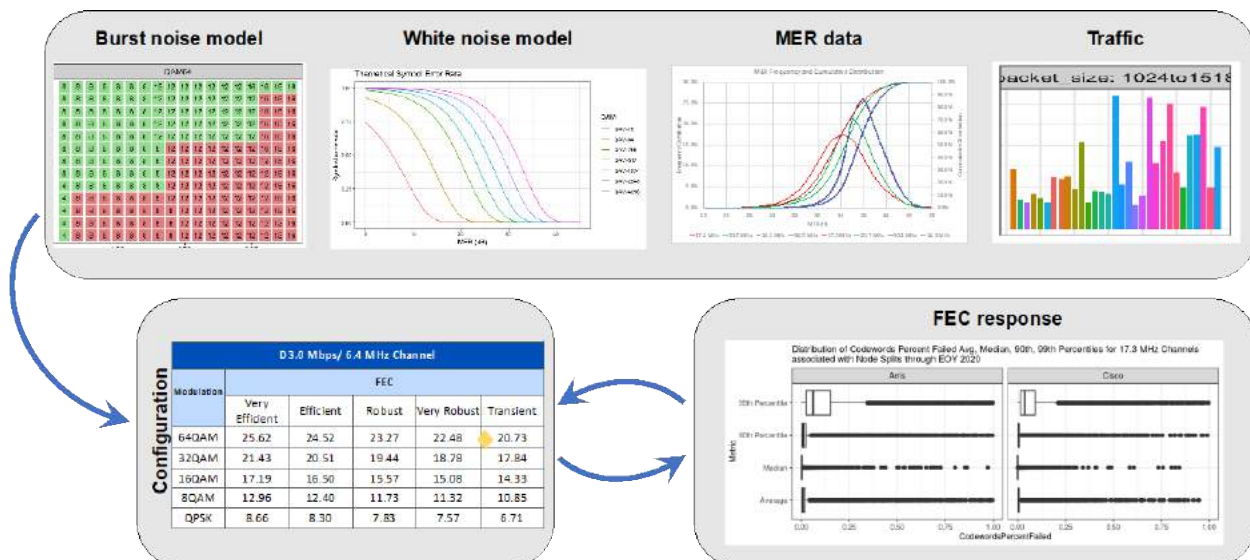


Figure 12 - D3.0 PMA System Comes Together.

4.3. Why D3.0 US Optimization increase network capacity?

The core idea behind the US profiles is based on changing the FEC configuration to play the trade-off between increasing codeword packing efficiency and increasing robustness against noise (be it white

noise or high frequency transient noise). The concept is illustrated in Figure 13, which shows an illustration of the anatomy of a FEC codeword for a range of modulation profiles. In this example, moving from the transient to the very efficient profile would allow transmitting the same packet while leaving sufficient time for two additional DNS requests to a website (such as “www.ieee.org”) that otherwise would have been occupied by un-necessary physical layer and time overhead. The codeword includes a preamble at the beginning, a guard interval at the end, and portions for data (payload) and parity bytes in between. The lengths of all of these components are defined as part of an US profile configuration. In addition, D3.0 allows the definition of different FEC configurations for different traffic packet types (short data grants, long data grants, and voice.) This is where the problem becomes an optimization problem: Knowledge of traffic and noise patterns can inform proper profile configuration, to achieve high packing efficiency (defined as data payload length/codeword length), while maintaining error rates at an acceptable level. Rather than constructing the “right” profile for each US interface, we created a global template of 32 profiles covering a range of operating regimes and QAM modulations (a subset of which is shown in Figure 11 for 6.4 MHz wide channels). Notice how varying the packing efficiency is achieved by changing the FEC configuration: the long data grant for profile 251 (the most efficient 64-QAM profile) has a payload length of 247 bytes and a parity length of 4 bytes, while profile 271 (designed to deal with transient noise) has a payload length of 121 bytes and a parity length of 16 bytes. Consequently, profile 251 has a speed that is approximately 20% larger than profile 271. Note that the values stated above are not arbitrary. They were designed in consideration of the actual traffic packet size distribution on the network, and the type of transient noise typically experienced due to DC power supplies (the so-called common mode disturbance (CMD) noise described in Ref. 7).

Minislot Timing Example 2: 741 Octet Mid-size Packet

64 QAM, 6.4 MHz, Very Efficient, Very Robust, Transient Noise

- Long Data Grant
 - protect stronger
 - more overhead
 - less efficient
- Very Efficient profile takes ~12% fewer minislots to transmit same data
- More minislots available for other packets
- Transient protects against transient Ex: power supply impulse noise

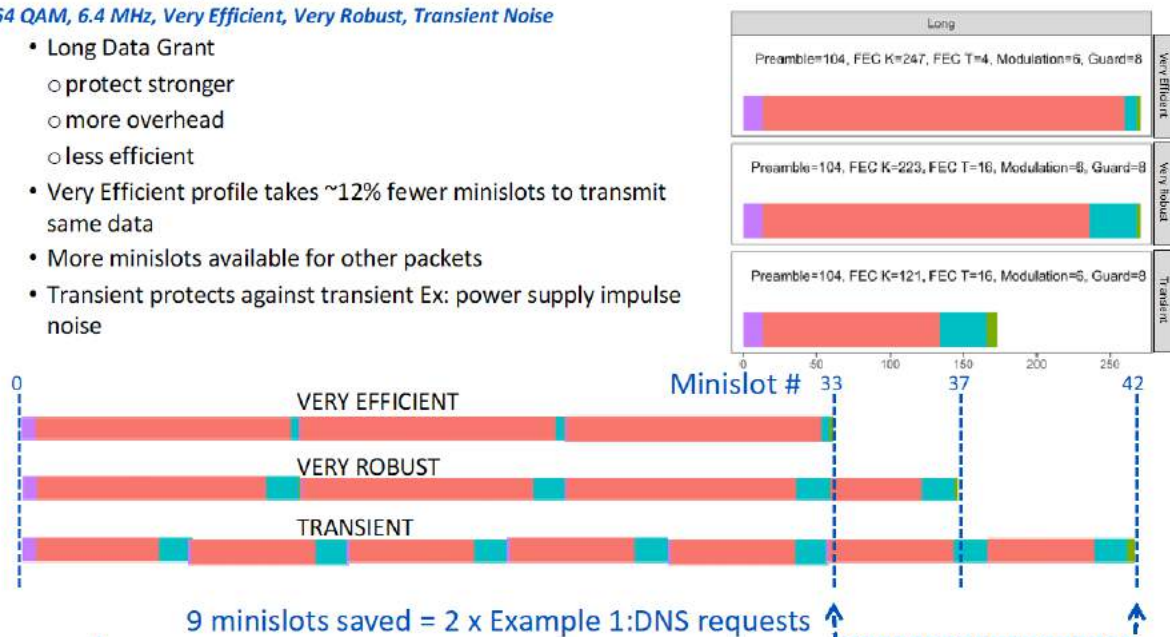


Figure 13 - Anatomy of a FEC codeword showing example of different codeword configurations for the D3.0 US long data grant.

4.4. The Reinforcement Learning approach to US PMA

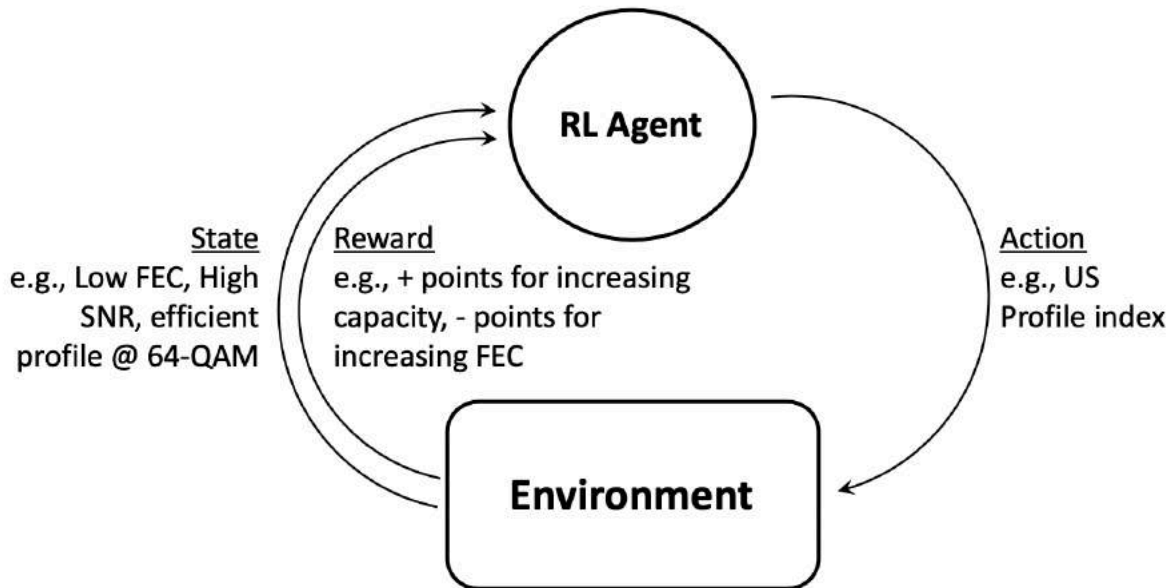


Figure 14 - Schematic of the Reinforcement Learning system for US PMA.

The general scheme for the Reinforcement Learning (RL) problem is shown in Figure 14. The goal of the agent, as it interacts with the environment, is to learn an optimal policy that maps states into actions. In the context of upstream PMA, states represent channel telemetry and configuration, while actions represent the upstream profile choice. The ideas behind RL are based on decades-old concepts in control theory and dynamic programming. What is unique to RL is the concept of learning from “trial and error.” That is, the agent will at times purposely select a random, non-optimal action in order to continue exploration of the environment. This is useful for two reasons: (1) The number of states, in combination with possible actions, is too vast to sample in a systematic way. (2) Even with sufficient knowledge of the dynamics of the environment at a given point in time, the environment is typically non-stationary. Changes in these dynamics warrant that the agent maintains a level of exploration over time. Our first attempt at representing the US PMA control problem as an RL problem included defining the following components:

- **States:** FEC error rates, traffic volume, and device SNR aggregated to the channel level are discretized by binning the continuous variables according to some given thresholds (e.g. uncorrectable error rate > 1% is High, <1% but >0.1% is Medium, and <0.1% is Low). The profile modulation (QPSK to 64-QAM) and profile type (most efficient to most robust) are discrete to begin with. The resulting state space includes ~6,000 states.

- **Actions:** We limit the agent’s actions to changing the profile level by +/-3 steps (in either direction, i.e. upgrading or downgrading by n steps), maintaining the same state, or transitioning in or out of the special profile designed for transient noise. Thus, the total number of possible actions is 9.

- **Rewards:** These are scalar values allocated to the agent following each action taken. They are designed to incentivize the agent to maximize returns in the long run. Rewards proportional to the change in profile level are awarded (e.g. upgrading the profile by 3 steps incurs a reward of +3). However, experiencing an uncorrectable error rate above 1% incurs a reward of -10, as this is deemed to be negatively impactful to

the customer experience.

- **Policy:** The policy is a mapping from states to actions—i.e. it represents the decision process. The agent starts with an initial policy that resembles the fixed rule-based approach in the live production PMA pipeline. Example: if FEC rates are below a certain required threshold and the SNR level is above another required threshold, then the channel is reconfigured with a profile that is one step higher than the current one.

With the above in place, the goal of the agent is to discover and maintain an optimal policy. This is done by maintaining a degree of exploration. Typically, the agent would select a random action 1% of the time and the optimal action 99% of the time, under what is known as the epsilon-greedy policy (with $\epsilon=0.01$). Initially, there will be a subset of channels running on a randomized profile between update cycles (currently 6 hours). The risk of causing an adverse effect is mitigated by the fact that the action space limits the magnitude of the upgrade to 3 steps. Furthermore, as more states are explored and convergence to an optimal policy is achieved, the amount of exploration could be dialed down to below 1%. Within the core RL algorithm, as new state-action pairs get explored, the agent collects its rewards based on changes in the environment (perceived through the state) and updates a table representing the long run return for each combination of state-action. This table can be trivially converted into a “most optimal” policy at the time. The update process itself is founded in theory relating to Markov Decision Processes (MDP). These are a special type of processes in which the current state is sufficient to describe all the preceding history of the process. Modeling the RL problem as an MDP allows the use of a host of sampling-based approaches to efficiently update the state-action values. In addition, MDP theory predicts convergence to an optimum policy, given a stationary environment and sufficient number of update iterations. Our current efforts are focused on implementing the RL algorithm and testing it on a small subset of production CMTSs, in order to assess how well it performs against the existing static policy.

5. Conclusion

As of this writing, Comcast is currently optimizing capacity across thousands of CMTSs and hundreds of thousands of OFDM and D3.0 upstream channels for more than 50 million DOCSIS gateways and cable modems. The system processes tens of terabytes of data and performs approximately 500,000 recommendations and transactions per day to optimize the upstream and downstream channels. To date, this has yielded capacity improvements of more than 30% in the downstream (towards customers) direction, and of approximately 20% in the upstream direction (from customers.)

To achieve scale and very high levels of reliability, PMA was developed as a serverless, elastic, cloud-native solution—taking advantage of distributed compute, storage, and network options. One of the primary design patterns was to allow for component independence: loose coupling with strong interoperability to promote feature velocity and a high cadence of delivery. Adopting these basic principles enabled each of the functional components and sub-components to have independent, automated, and continuous integration and deployment trains. This allowed for system updates to scale to business needs, often being updated seamlessly and in parallel several times a day.

There are real opportunities to capture additional efficiencies from invested capital in the network. Along this journey we learned a couple things:

- *Keep it simple: It matters to scale.* Network capacity is a complex environment. Modern access networks are complicated; DOCSIS technology is complicated; CPE devices are complicated and required many firmware updates. CMTS solutions are heterogeneous and required many new features to be developed. The number of devices and variety of devices are complicated; their

different firmware versions, also complicated. Complex systems often exhibit emergent and unexpected behavior and require special consideration when building to scale. Best to keep it as simple as possible (advice perhaps easier said than done!)

- *Leverage the cloud.* The benefits of the cloud are real, and public cloud is an accelerator. Leveraging public cloud resources allows you to focus resources on key problems. It allows you to reliably reach scale, at a velocity that meets business demands.
- *Keep learning through iterative builds.* For best results, build on platforms that can be updated seamlessly and constantly tested. We could not have anticipated all the conditions that we encountered and benefited greatly from a team and culture that was built on iterative learning.
- *Be data-driven by measuring what matters.* Along the entire journey, from early analysis to production dashboards and quality metrics, trusting the data and developing a data-driven strategy allowed us to match practice to theory and build confidence. Being data driven will drive you toward success.

Abbreviations

AE	Analytics Engine
AI	artificial intelligence
API	application programming interface
AWGN	additive white gaussian noise
CAGR	compounded annual growth rates
CLI	command-line interface
CMTS	cable modem termination system
CMD	common mode disturbance
DOCSIS	Data Over Cable Service Interface Specification
D3.0	DOCSIS 3.0
D3.1	DOCSIS 3.1
DE	Data Engine
DS	downstream
FDX	Full duplex
FEC	forward error correction
HFC	hybrid fiber-coaxial
JSON	JavaScript Object Notation
LDPC	low density parity check
MER	modulation error ratio
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
PHY	physical layer
PLC	physical link channel
PMA	Profile Management Application
QAM	quadrature amplitude modulation
RL	Reinforcement learning
SCTE	Society of Cable Telecommunications Engineers
SNR	signal to noise ratio
UGS IUC	unsolicited grant service interval usage code
US	upstream

Bibliography & References

1. “A Machine Learning Pipeline for D3.1 Profile Management”, M. Harb, J. Ferreira, D. Rice, B. Santangelo, and R. Spanbauer, NCTA technical paper, 2019.
2. “Our Network”, Comcast corporate website, <https://corporate.comcast.com/our-network>
3. “Comcast Network Investment”, Comcast press update, <https://update.comcast.com/download/15781/>
4. “As remote work exploded, Comcast turned to AI to keep the internet running”, FastCompany article, <https://www.fastcompany.com/90519167/as-remote-work-exploded-comcast-turned-to-ai-to-keep-the-internet-running>, 2020.
5. COVID-19: How Cable's Internet Networks Are Performing: METRICS, TRENDS & OBSERVATIONS, <https://www.ncta.com/COVIDdashboard>.
6. <https://www.nctatechnicalpapers.com/Paper/2002/2002-optimizing-the-last-mile>
7. Convolutional Neural Networks for Proactive Network Management: Developing Machine Learning Models to Detect and Classify Impairments in D3.1 OFDM Channels, SCTE Cable-Tec Expo 2020, Ferreira, Harb, Subramanya
8. “DOCSIS 3.1 Downstream Early Lessons Learned”, J. J. Downey, NCTA technical paper, 2017.
9. What Gets Measured Gets Done / What Gets Analyzed Gets Transformed, Mutalik, Rice, Subramanya, Wang, SCTE IBSE Cable-Tec Expo 2018, <https://www.nctatechnicalpapers.com/Paper/2018/2018-analytics-for-a-wider-deeper-network-view>

Additional References

ANSI C63.5-2006: American National Standard Electromagnetic Compatibility–Radiated Emission Measurements in Electromagnetic Interference (EMI) Control–Calibration of Antennas (9 kHz to 40 GHz); Institute of Electrical and Electronics Engineers

The ARRL Antenna Book, 20th Ed.; American Radio Relay League

Code of Federal Regulations, Title 47, Part 76

Reflections: Transmission Lines and Antennas, M. Walter Maxwell; American Radio Relay League

Operationalizing Streaming Telemetry and Machine Learning Model Serving

Customer Experience Automation

An Operational Practice prepared for SCTE•ISBE by

Nick Pinckernell

Distinguished Engineer

Comcast

183 Inverness Dr W, Englewood CO 80112

nicholas_pinckernell@comcast.com

Scott Rome

Principle Researcher

Comcast

1800 Arch St, Philadelphia, PA 19103

scott_rome@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. History	4
3. Customer Experience as it relates to ML	5
3.1. Use Case	5
4. Telemetry.....	5
4.1. Attributes.....	6
4.2. Usefulness.....	6
5. Platform	7
5.1. General Architecture.....	7
5.2. Implementation.....	9
5.2.1. Technology choice	9
5.2.2. Architectural differences.....	9
5.2.3. Message bus performance.....	10
5.2.4. Container orchestration.....	11
5.3. Feature engineering performance and guidelines.....	12
5.3.1. Min-Max normalization	12
5.3.2. Sum or counts.....	13
5.3.3. Mean	15
5.3.4. Other operations	15
6. Model Serving.....	16
6.1. Requirements	16
6.2. End-to-end flow	16
7. Scale	17
8. Model.....	17
8.1. Feature Engineering	19
9. Results	20
10. Looking Ahead	20
11. Conclusion	21
Abbreviations.....	21
Bibliography & References	22

List of Figures

Title	Page Number
Figure 1 - Telemetry collection flow.....	7
Figure 2 - Feature engineering and model invocation flow.....	8
Figure 3 - End-to-end telemetry and ML processing flow.....	8
Figure 4 – Message bus parallelism.....	10
Figure 5 – Message bus consumer scale matching	11
Figure 6 – Container parallelism	12
Figure 7 – Container isolation	12
Figure 8 – Inference graph example.....	16
Figure 9 – End-to-end platform flow	17
Figure 10 – Model flow	18
Figure 11 – Model architecture	19

List of Tables

Title	Page Number
Table 1 – Sample telemetry	10
Table 2 – Message bus consumer scaling	11

1. Introduction

Customer Experience, Telco, Machine Learning, Automation... all sounds a bit dry, no? If not, this should help shed light on a specific use case, but also provide details and lessons learned while building the solution. If so, read on to perhaps find some surprising details. Many are aware of the impact from Machine Learning but may not be aware just how many portions of the enterprise it is now altering -- or the magnitude of those changes.

With building an ML platform in mind, specifically to improve and automate the customer experience, this paper will illustrate how we accomplished this. It will also detail lessons learned and insights not only on the data itself, but on the architecture and thinking behind it.

Another focus is the challenges and differences of scaling and maintaining the Machine Learning components of the architecture. The highlight here is that while Proactive Network Management (PNM) [1] has laid out a number of excellent methods on how to collect and analyze network telemetry from CPE and other headend equipment, it has not covered the aspects of what to do or how to handle the data with respect to utilizing ML models.

This paper will begin with the Customer Experience use case, and work from that point through to the end solution. A number of open source technologies are referenced as possible implementations for components. Other components can certainly be used.

The desire to share this information stems from the fact that some of these solutions are not easy. They require not only multiple resources to develop the front-end user interface, but the back-end platform as well. After the ML models have been trained, the task of building the platform, scaling, testing specific tools and gluing everything together is still rather manual and prone to performance and optimization challenges. With those things in mind, the details throughout this paper should aid the reader in determining directions to go when considering, building and scaling such a system.

2. History

Telemetry collection and polling have been widely covered in the industry with both positives and negatives. The positives almost always outweigh the negatives, in terms of providing the ability to determine outages quickly, gain insights on which partitions and elements in the network require maintenance and many other useful applications. Using telemetry with machine learning is also fairly old, just not commonly referred to as “machine learning”. Take adaptive equalization in digital communications as an example. The parallels are many between parts of adaptive equalization and how machine learning algorithms work. Or, take the DOCSIS upstream pre-equalization process as another example. For each burst of data in a transmission, the preamble is used just like a training set for a supervised ML algorithm, but for the CMTS’s upstream adaptive equalizer. Likewise for the downstream, with a blind-equalizer -- it is the same as an unsupervised ML algorithm. Overall, there are inputs, derived coefficients and an optimization problem – which is to minimize the mean square error of the outputs.

Applying ML algorithms to other telemetry angles is simply considering different shapes. With polling or collection systems, gathering telemetry from CPE, headend or data center equipment, and funneling that into distributed computing environments to do essentially the same thing the CMTS was doing with adaptive equalization, is a larger scale scenario with potentially more complex models. This method of analyzing telemetry is just the next evolution, after analyzing telemetry for outages or full spectrum data

for various distortions -- but now using more complicated methods such as ML models to tease out new artifacts and correlations in determining network health.

3. Customer Experience as it relates to ML

What is customer experience? According to customer experience futurist Blake Morgan, it “really boils down to the perception the customer has of your brand” [2]. In the use case laid out, it is the perception of the self-service customer experience.

3.1. Use Case

The outcome of the project is to automate simple common customer service questions and return immediately useful feedback or solutions to the posited issue or question. The goal here was two-fold: one, to improve the customer experience by knowing their issue before they do and if possible remedy it, and two, to reduce the number of calls to customer service.

When left open-ended, the outcome is intractable, given the large problem scope of all possible issues customers may face while going through all the various customer service avenues available to them. Therefore the scope and requirements of the outcome were reduced to answer a few basic questions, then continue to build out the capability to handle more queries from there. These initial questions were:

- Is there an issue with my internet service?
- Is there an issue with my video service?
- I have a billing or account inquiry?

All of the problems require use of machine learning. The first two require telemetry data and additional machine learning as well. For the third question, the goal was to handle simple questions initially, such as “why is my bill different than last month?” or “how much does my TV package cost?” and has since been expanded to handle more questions.

With the goals and outcome clearly defined, the solution required thinking on how to accomplish this, given the large number of telemetry and other data sources available.

4. Telemetry

The breadth of telemetry for most telecommunications companies and network operators is vast. Data comes from a number of sources which generally are customer devices, headend or data center equipment or public devices. Customer devices or CPE (Customer Premises Equipment) are usually comprised of cable modems, Wi-Fi Access points, cable boxes and the ever growing list of IoT devices. Central locations such as the headend, Point of Presence (POP) and data centers tend to be the geographic source for equipment such as CMTSs, access networks, video sources and more.

The correlation between our initial questions and data sources were fairly straightforward:

- Internet service problem: telemetry from the cable modems and CMTSs
- Video service problem: errors from our Video Backend Service (VBS)
- Billing / account inquiry: billing data system

During the exploration of datasets, a single dataset was found that provided attributes for both the internet and video service questions, which was errors from the OS that runs on both cable modems and STBs.

The informational and error messages from devices built on the Reference Design Kit (RDK) [3] turned out to be very valuable, as both the Set-top Box (STB) and Cable Modem (CM) use the same transmission medium from the premise to the headend or hub site.

4.1. Attributes

Telemetry from the RDK consists of a number of attributes useful to determine if there are issues from the CM side, STB side or both. It contains general information about the CPE:

- CPU utilization
- Memory utilization
- System load
- Wi-Fi signal strength
- Etc.

Perhaps the most useful data from RDK, however, are error codes and counts. These errors relate to dropped packets, firmware errors, signal errors, etc. The specific attributes which were the most impactful as part of the machine learning model were “rf_error_router_ip_loss”, “rf_error_ipv6pingfailed”, and “rebootreason=unknown”. These attributes increase the accuracy of the models in determining whether or not the issue is related to internet service / high speed data (HSD).

While the RDK telemetry provided beneficial information for both HSD and video, the VBS data provided the majority of telemetry needed to determine if the customer issue was video-related. The VBS data contains such telemetry as:

- User interactions with their TV via the remote
- Errors displayed on-screen to the customer
- Errors with satisfying a request to the customer
- Etc.

For the billing questions, having data about the customer billing history, specific elements on the bill and differences between those elements over time provided the majority of features needed to satisfy the initial use-case.

4.2. Usefulness

In “Observing home wireless experience through wifi APs” [14], it was shown that many in-home Wi-Fi markers were able to characterize wireless experience. In our work, the definition of telemetry includes error and system logs produced by RDK, which is unique to our use case compared to the literature. We will detail a few examples in this section to give the reader a picture of their utility.

The telemetry includes readings of Wi-Fi signal strength (known as RSSI, for Received Signal Strength Indicator) and channel utilization. RSSI is a negative value (-100,0) and a value below -80 is considered poor signal strength. Therefore, it is clear that this telemetry feature can be useful in identifying devices that have an impacted customer experience. Likewise, channel utilization data can indicate when a given Wi-Fi channel frequency is saturated from too many devices or too much traffic. High utilization is an indication that one may need to change the Wi-Fi channel the router is using for signal transmission.

The RDK logs are another source of events which may impact a customer’s service. For example, if a reboot occurs, there are multiple keys that indicate the event has taken place, and in some cases, why the reboot occurred. Unscheduled reboots are generally strong indicators of a customer experience issue.

Moreover, there are system logs when processes on the box restart, which can also impact service, depending on the process. (Such processes are programs running on the thin version of Linux on the gateways). There are other markers like “system_uptime”, which gives the time since a last reboot, and “rf_error_ipv6pingfailed”, which indicates a count of pings to a fixed IPv6 server address that have failed since the last log. In total, there are over 750 unique keys which appear in RDK and more are added over each release.

5. Platform

5.1. General Architecture

All of the features learned from exploratory data analysis need to be processed, parsed and transformed so they can be passed into the model. During offline training of the model, the researchers perform these tasks but usually at smaller scales. Or, if full-scale, it usually doesn’t run constantly. Also in the research environments, the lack of introspection, scaling (depending on the researchers’ environment), monitoring and alerting is unacceptable for running a supported product. These are among the goals of an ML platform.

The basic components of the platform, in order of data flow, are:

- Data producer
- Raw data consumer
- Feature engineering
- Model invocation
- Inference storage or action

This typically involves one or more teams and one or more platforms to handle the data. It is common for the platform responsible for polling or collecting data from the CPE to be handled by one team, while a different team is responsible for the ML platform.

When greatly simplified, data collection can be boiled down as seen in Figure 1. Here, some process (usually many parallel processes) is responsible for connecting or listening to the CPE and other devices to collect, format and aggregate telemetry. JSON is typically used for data formatting, though other formats are becoming more common, such as Protocol Buffers and Apache Arrow objects. After the data is formatted, the telemetry collection system will publish the messages to a message bus for consumption from other teams within the organization. For our systems, we use Apache Kafka, as it is stable and scales well.




Figure 1 - Telemetry collection flow

Once the data has been published to the message bus, it is decoupled from the data collection platform and available to be used by the ML platform. A simple flow for processing the telemetry events from the message bus is detailed in Figure 2. Here a message bus consumer process will listen to the stream for new messages. Next the message will be parsed, and various transformation and/or normalization steps will be performed on the data to prepare the dataset to be passed into the ML model. Before or after parsing and feature engineering, it is common to store features into a database so they can be later retrieved. This is useful when building aggregate feature sets (such as time windows, feature enrichment, etc.) so the model may be passed a richer and potentially more useful set of features. Once the features are ready, the model is invoked, which produces an inference or prediction. This output is then handled by possibly storing that in another database, distributed filesystem or even published to another message bus to do something with that model output.




Figure 2 - Feature engineering and model invocation flow

While these two systems are usually logically separated and loosely coupled, thanks to the message bus architecture, they do need to combine as shown in Figure 3. Here the full picture is seen with both the telemetry collection platform and the machine learning platform receiving telemetry.

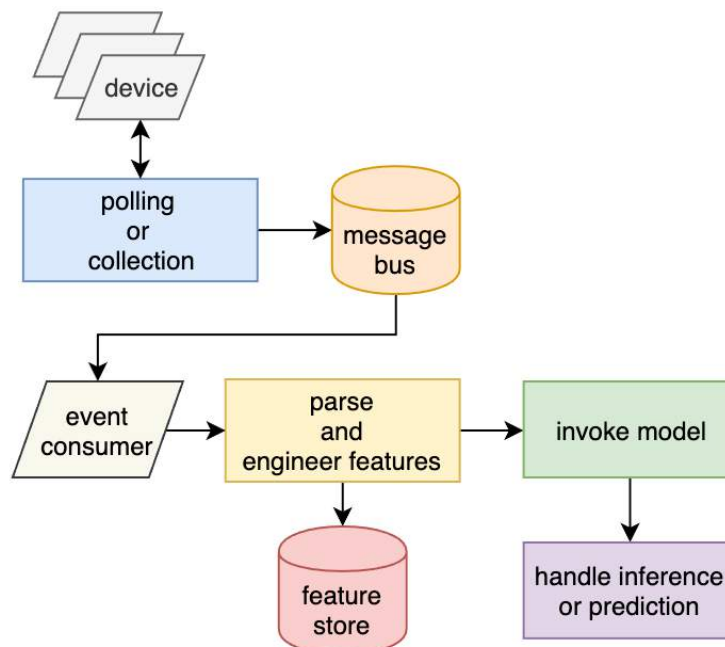


Figure 3 - End-to-end telemetry and ML processing flow

5.2. Implementation

5.2.1. *Technology choice*

The end architecture contains the same components and flow as shown in Figure 3, but there are differences in how those components look after scaling and with specific technology choices. For the flexibility to run programs of any type, Docker was chosen to contain some of these complex applications. For scalability, a container orchestration layer is required to dynamically scale resources to match the needs of the incoming telemetry – in our case, Kubernetes was chosen here. For a scalable message bus, the choice was Kafka, as previously mentioned. Depending on the rate of operations and type of data in the features, a range of database technologies and types can satisfy that requirement. For example, if you have relational data, a more traditional SQL database such as MySQL or PostgreSQL might be chosen. For a NoSQL database, perhaps Apache Cassandra or MongoDB. For a NoSQL cache layer, CouchDB, Redis, or if SQL, then MemSQL. This is by no means a comprehensive list of database technologies but do represent a few options to investigate, should you have a specific set of requirements for your feature store. For our solution, Redis was chosen primarily for performance reasons and data retention. The research team determined that the largest useful window of data was 24 hours and older historical data beyond that was less impactful. It is also relatively quick to repopulate, as well, in the case of catastrophic cluster issues.

Kubeflow was chosen for model serving not only because it works well with Kubernetes but contains a number of required features as well. More specifically, a sub-component of Kubeflow called Seldon Core was chosen due to its flexible inference graphs, monitoring and A/B testing capabilities. Finally, Python is the language of choice for deploying models due to the fact that the models will require little or no code changes to be deployed. Also, many Python specific packages such as NumPy are very performant and don't have great equivalents in other languages. This also allows the model owner or team less familiar with production operations to maintain and re-deploy their model as needed. Other scientific languages, such as R, have far less library and performance support for writing and running general components.

5.2.2. *Architectural differences*

The differences between our general architecture in Figure 3 change when we consider Kafka and Kubernetes. Starting with Kafka, it is important to understand a bit about how Kafka works so that we can scale appropriately. Without getting too deep into Kafka's architecture, there are three basic components: the topic (which consumers and producers use), the broker (a node in the Kafka cluster that handles topics) and partitions (a sharded piece of the topic). When we have high volume topics, those usually need to be partitioned out so we can scale throughput horizontally. In Figure 4, the telemetry polling architecture would publish messages to a specific Kafka topic with b_n brokers and P_m partitions.

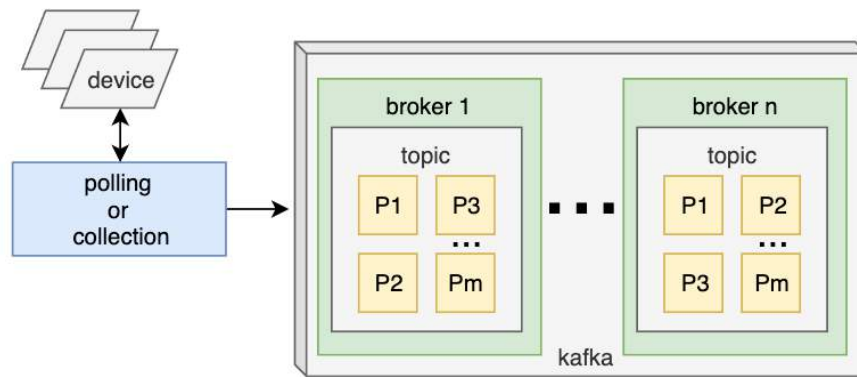


Figure 4 – Message bus parallelism

To calculate n and m , the volume and message size need to be determined. For this example X, Y and Z are just example features. Actual features used in the system are Signal-to-Noise Ratio (SNR), RSSI and others. Take the following row as an example:

Table 1 – Sample telemetry

MAC address	Customer ID	X	Y	Z
aa:bb:cc:dd:ee:fa	123456780	2.156	True	3
aa:bb:cc:dd:ee:fb	123456781	3.211	False	7

If each of these records were represented as individual messages, then they might look like this in JSON.
`{"mac": "aa:bb:cc:dd:ee:fa", "customerid": "123456780", "x": 2.156, "y": true, "z": 3}`

However, the data is usually far from optimal, introducing additional parsing overhead or unnecessary extra data. Such as an unnecessary element as shown below “pollresult” or perhaps elements represented as strings instead of their actual datatype.

```
{"pollresult":
  {
    "mac": "aa:bb:cc:dd:ee:fa", "customerid": "123456780", "x": "2.156", "y":
    "True", "z": 3
  }
}
```

5.2.3. Message bus performance

Using this less than perfect data, we can now calculate the volume and size of the messages to determine the correct number of brokers and partitions for the Kafka topic. Using our system as a more realistic example for the message size, the mean message size of 1.4 kilobytes is used. Assuming there are ten million devices being polled every minute and the JSON is raw (uncompressed) on average 1.4kB, and the message above, the throughput would be $\frac{1e7 * 1400}{1000^2} * \frac{1}{60} = 233. \bar{3} / \text{MBs}$. Using 50MB/s per broker as a rule from Dropbox’s Kafka Throughput limit post [9], rounding up we would require 5 brokers at 50MB/s each, giving us a reasonable 250MB/s throughput. To determine the number of partitions, multiply the number of brokers by 10, which yields 5 brokers with a single topic of 50 partitions. Of course, this is just an example, and situations are different based on hardware type, network throughput and many other factors.




Figure 5 – Message bus consumer scale matching

The reason partition counts per topic are important relates to the performance gain of matching consumer processes or as a multiple of the number of partitions. Here, with 50 partitions, you may need to scale the message bus consumers as shown in Table 2.

Table 2 – Message bus consumer scaling

Consumer count	Multiple	Comment
5	10	If message processing is very fast, single consumer may be able to handle multiple partitions each.
25	2	Many more consumers handling partitions. If message processing is more computationally intensive, each consumer might only handle two partitions.
50	1	A single consumer per partition is required if the code is not only parsing, but performing some action or invoking a ML model inline as shown in Figure 5.

Final thoughts on consumers and messaging or streaming systems: There are many combinations of great open source projects which can be used instead of writing the consumer logic manually. For example, systems like Apache Flink and Apache Beam, as well as offerings from the various cloud providers, can perform many of the functions described here.

5.2.4. Container orchestration

Once the number of consumers has been determined, a scalable, fault-tolerant environment is needed to run the consumers and models. The flexibility of the runtime environment is crucial to operating and maintaining ML models and pipelines, which made Kubernetes and Docker a natural choice. Kubernetes also allows elastic horizontal scaling, to scale up for spikes and scale down for dips in processing.

Zooming in on the general architecture in Figure 3, the message consumption and feature handling, as well as the model invocation components -- all translate to Kubernetes pods. Once the engineer or researcher creates the model, it is pushed to a container repository, then pulled down in Kubernetes to run and monitor, as shown in Figure 6, until there is a change. Kubernetes can scale up and down by adding and removing replicas of a particular pod corresponding to the addition and removal of worker nodes.




Figure 6 – Container parallelism

A benefit of Kubernetes is that pods of any type can be run together on the same cluster. If the consumer pods for example are written in Java, and the model code is written in Python, then both can run simultaneously on the same worker nodes, thanks to Docker, and shown in Figure 7.




Figure 7 – Container isolation

5.3. Feature engineering performance and guidelines

Just before the model can be invoked, the raw features from the polling architecture require parsing and some manipulation specific to the ML model. Many types of manipulations exist, such as min-max normalization and standardization. For specific ML domains such as Natural Language Processing (NLP) or Deep Neural Networks (DNN), the data may need to be tokenized or one-hot encoded. Three of these normalizations are common, so it is necessary to look at them with the corresponding features and how those features get changed.

5.3.1. Min-Max normalization

To rescale features between 0 and 1, the formula is $x' = \frac{x - \min(x)}{\max(x) - \min(x)}$. However, this requires that the maximum and minimum values of the features are available. Many systems may only contain partial feature sets from which the minimum and maximum value cannot be derived. Note that the term “feature set” is just a series or list of features, and order may be important. Also common is to scale features to a particular range $[a, b]$. For this, the formula is slightly different $x' = a + \frac{(x - \min(x))(b - a)}{\max(x) - \min(x)}$. If the features

require these normalizations, the minimum and maximum may be provided, or the telemetry will need to be collected for an acceptable period of time to derive acceptable values.

5.3.2. Sum or counts

Adding or counting features for a particular time window or sample is a common normalization. However, the computational complexity when performing feature engineering at scale here can be deceiving – especially when referring back to the example of ten million devices polled every minute. Take a hypothetical situation where the desired prediction is a usage pattern spike given two features, bytes in and bytes out. Assume that research has found 30 minute windows to be optimal for predicting if there will be a spike. The total feature size for all devices would be $1e7 * 30 = 300e6$ or 300 million. If the features were simply a MAC address, bytes in, bytes out, and timestamp, the JSON representation might look like

```
{ "mac": "aa:bb:cc:dd:ee:fa", "bytesin": 1287630, "bytesout": 58360, "timestamp":  
  "2020-07-01 02:31:05" }
```

The feature set size for all ten million MAC addresses would be $\frac{\sim 103 \text{ bytes}}{1024^2} * 10e6 \text{ devices} * 30 \text{ minutes}$. While this is only ~29.3GB worth of features which is relatively small, the compute time may be large. Consider another example: Python code as pseudo-code. There are a number of details left out, such as actual timestamp calculation, and data structure details:

```
# iterate through current polled messages from the kafka topic for each device  
  
for features in current_features_from_kafka:  
  
    # get history  
  
    history = get_history_for_mac(features["mac"])  
  
    history.append(features)  
  
    # iterate through history  
  
    summed_features = {"bytesin": 0, "bytesout": 0}  
  
    for history_features in history:  
  
        # expire old items  
  
        if history_features["timestamp"] > thirty_minutes_ago:  
  
            # iterate through features
```

```

summed_features['bytesin'] += history_features['bytesin']

summed_features['bytesout'] += history_features['bytesout']

```

Say there are 50,000 messages from Kafka, each having 30 minutes of history, with the computational complexity of $O(n^2)$ (where n is the number of steps), which results in $50000 * 30 = 1,500,000$ iterations. If all ten million are polled every minute, that is 300 million iterations per minute. All of this requires more replicas to scale or, some simple tuning of the algorithm. If the math is commutative, it can be parallelized out of order, and for a simple count, it is. However, here the values may contain negatives, which means it is non-commutative and must be executed in order, serially. However, these can certainly be parallelized as the only requirement is that each device needs its feature sets to be computed in order. By changing the algorithm to remove iterating through the entire history set for each new message, we can reduce the second loop from 30 iterations to just two operations when storing the summed features as well in a database. There are obviously a few more steps here, such as the back and forth from the database and type conversion (if necessary), but it can all add up -- which can, in some cases, dramatically increase costs. With the improvement, for the 50,000 messages from Kafka we're now only doing 50,000 iterations instead of 1,500,000 and a time complexity of $O(n)$.

```

# iterate through current polled messages from the kafka topic for each device

for features in current_features_from_kafka:

    # get history from feature store

    history = get_history_for_mac(features['mac'])

    summed_features = get_summed_features_for_mac(features['mac'])

    # remove the first (oldest) and decrement from summed_features

    oldest_history = get_expired_feature_from_history(history)

    summed_features['bytesin'] -= oldest_history['bytesin']

    summed_features['bytesout'] -= oldest_history['bytesout']

    # add the current features

    summed_features['bytesin'] += features['bytesin']

    summed_features['bytesout'] += features['bytesout']

```

Said another way, we can look at it as a sum of series, where those series may contain negatives. Let n equal the number of historical feature sets for the series a , and m equal the number of features within each historical feature set a_i . The total sum S_m is the series of all historical feature sets:

$$S_m = \sum_{i=0}^{n-1} \sum_{j=0}^{m-1} a_{ij}$$

Now let c equal the current message being processed with p features. Then the historical features sum the current features:

$$S_m = \sum_{k=0}^{p-1} S_k + c_k$$

If the previous sum of the historical feature sets is known, then let h equal that series, and the original operation can be rewritten more efficiently as a single sum. This does not take into account the subtraction or decrementing of historical feature sets that expired from the desired time window (30 minutes for this example):

$$S_m = \sum_{k=0}^{p-1} h_k + c_k$$

Keep in mind that this assumes the series and elements are the same between series. In real-world data, the elements for each piece of telemetry may be different, resulting in sparse data structures which require different thinking to process efficiently.

5.3.3. Mean

The mean here is a bit more complicated because without iterating through the entire history an actual mean $\frac{1}{n} \sum_{i=1}^n a_i$ can't be derived. To get around this and get a "good enough" mean, we turn to a simple moving average, so we don't have to iterate through the entire historical time window of features. With a simple moving average $\frac{1}{n} \sum_{i=0}^{n-1} p_{M-i}$ is used as we have the historical features in the database and the code is similar to the sum example.

5.3.4. Other operations

For operations such as standard deviation, minimum or maximum values are more complicated to process given a single value. There are potential ways to calculate these things given more values, such as minimum or maximum, if you know the top or bottom 5 for a particular feature. This requires keeping those lists and can be cumbersome -- which may defeat the purpose of computational efficiencies, unless the historical time window or number of messages therein is large.

6. Model Serving

The last component of the system is the model, and how it gets exposed to users or other parts of the organization. This is perhaps the most straightforward piece of the system, not considering what happens to the output (prediction or inference) from the model. The simplicity comes from a project called Seldon Core which allows almost any type of ML model to be invoked with Python and served from Kubernetes. In Figure 7, not only is the model potentially running alongside the feature engineering on the same node, but in separate containers, it is also able to scale horizontally by simply changing the number of replicas. This can also be performed automatically with elastic scaling. All the major cloud providers support some kind of model serving and each has their own benefits and drawbacks depending on the situation. Running your own Kubernetes allows you to also run Prometheus or tie it into an existing monitoring solution within the organization, such as to leverage metrics and alarming for the operations support side of these platforms.

6.1. Requirements

Given the large number of components in the model, a system which supported core requirements such as multi-armed bandits, A/B testing and advanced inference graphs was required. Seldon Core was chosen as it satisfied all of these requirements and integrated well with Kubernetes and production monitoring systems. Seldon can support a user-defined inference graph, which allows custom configuration and combinations of models and components. It also allows for percentages of traffic to be handled by each, thus allowing for more complicated implementations such as A/B testing. As part of the user request flow to the model, it was important to ensure that customer HTTP sessions are sticky to a particular A/B model. In order to accomplish this, a custom Seldon router, which takes a seed value can be written to route requests to the same components in the inference graph depending on values in the payload (such as a customer identifier).

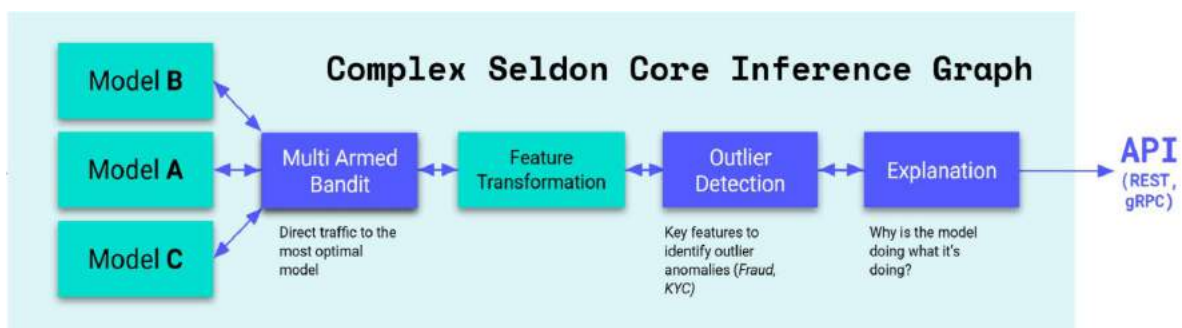


Figure 8 – Inference graph example

6.2. End-to-end flow

While the entire flow is shown in Figure 3, there is a critical piece missing. Figure 9 depicts the end-to-end flow, including customer interaction with the system to request predictions from the model.




Figure 9 – End-to-end platform flow

In the second flow with the thick blue lines, the customer interacts with the customer service application on their device or web browser. This invokes the model service and returns the prediction back to the application, then the user, about whether or not it believes they are experiencing an issue based on the telemetry and output from the application. If there is an issue, it attempts to specify a recommended fix action, or, as a last resort, refer the customer to customer service. This system has been successful at handling millions of customer requests quickly and accurately, improving the customer experience.

7. Scale

In our customer experience ML platform, we're generally handling 100k+ messages per second and executing 300-500k+ operations per second to our Redis clusters. As previously mentioned, this is accomplished with horizontal scaling from Kubernetes. In many of our platforms, we're handling more than just the RDK datasets, as our models require engineered features from many varying datasets to produce the unique outputs to improve the customer experience.

Our Kubernetes clusters generally run with 10 or more nodes at 16 CPU cores per node and 128GB RAM. The Redis cluster nodes have far fewer CPU cores (4) as the Redis process is single-threaded, yet they have much more memory, generally 256GB RAM. Our Redis nodes as well as the Kafka broker nodes have secondary 1TB drives for disk-based persistence, replication and retention.

8. Model

Our goal is to prompt a user to troubleshoot one of their services in order to increase engagement and troubleshooting inside the chat bot. Thus, our data is generated based on both customer feedback and the decision to show a prompt to a user. In this way, the model will impact the data generating process, which is different from standard classification or regression use cases, where the data generation is assumed to

be independent of the model’s output. For example, in image classification, if a model incorrectly labels a picture of a cat as a dog, there is no impact to the ground-truth label—that the photo is of a cat. In our case, we will have feedback based on the action we select, but we have no information on if the model took a different action. Our problem can be formalized in the framework of contextual bandits.

Contextual bandits algorithms extend the traditional multi-armed bandits problems by giving the agent a state to aid in decision making. Each data point can be thought of as a “round” of a game, where the state is the features of the model (“agent/policy”), the action is chosen by the agent, and the reward is the label. Contextual bandits problems are a special case of Reinforcement Learning, where the length of a round is 1. The goal of the agent is the learn to play the game in order to maximize the expected reward when following the agent’s action recommendations.

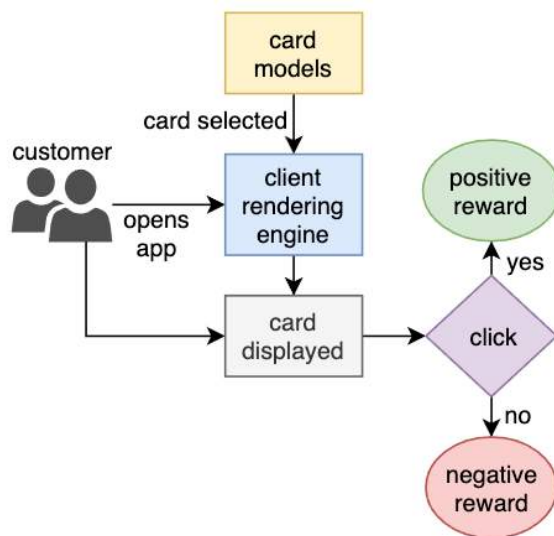


Figure 10 – Model flow

One drawback to reinforcement learning and contextual bandits in industry is that the algorithms are traditionally trained online. In the case of reinforcement learning, agents typically have access to a simulator of the environment (i.e. a game) and they learn as they play (c.f. “Playing Atari with Deep Reinforcement Learning” [17]). For contextual bandits algorithms, they often need to be deployed into production and learn as they serve customers [15], [16]. This is because the data distribution changes as the agent learns, as the agent will select new actions based on new information. This is different from the static training set found in classic supervised learning. Because of this, a suite of theoretical techniques has been developed to evaluate and train agents offline, namely “off-policy” techniques [18], [19]. Off-policy techniques have been employed in a variety of industrial applications successfully, including [20], [21]. This approach utilizes data from a production logging policy with sufficient randomization for training and evaluation of offline policies.

One of the large benefits of off-policy evaluation is the ability to estimate the true online performance of a contextual bandits model without deploying it into an A/B test. This allows the researcher to rapidly prototype many candidate models and choose the best model for an A/B test which impacts customers.

With that said, our model is a linear model which predicts the reward for each action (prompt) that we may show the user given the state, i.e. telemetry features defined above. After predicting the expected

reward for each action, we take the action with the highest reward and output that action to the production system. We train the model using various off-policy approaches [22] and select the model which performs the best on the off-policy evaluations of our metrics for production. Crafting reward functions is an art more than a science, and generally are designed by experts to maximize not only the business metric of interest but also downstream metrics. In our case, our reward function is a proxy for the click-through rate of the prompt, where we have a negative reward if a card was shown and not clicked. We scaled reward functions to be between $[-1,1]$, as it makes the training process more stable through smaller gradient updates.

All analysis was done in Jupyter Notebooks, including training. Models were defined and trained using Keras, a popular abstraction of neural network operations that sits on top of libraries like Tensorflow. In production, the model is invoked to select a card for the user upon application start.




Figure 11 – Model architecture

8.1. Feature Engineering

The telemetry we utilize comes in several forms: snapshots of physical characteristics of the Wi-Fi signal; device performance indicators usually expressed as a percent; and application logs which count the occurrences of different system messages, including error codes. We pass each of these types of features through a different pipeline to create appropriate features. All of our features are aggregations over window of time t , where it depends on the feature type.

For continuous features, we utilize statistical aggregations over the time window. Many of the continuous features are related to devices connected to the gateway, and so we will have multiple devices for a given customer. We take an aggregation of statistics calculated on each device as our final feature. For example, consider the RSSI: we calculate the mean, median, the standard deviation, and other statistics for each device, and then we aggregate these statistics again for a single value for each customer account. Time aggregations are used so that the model has is able to learn a sense of what’s “normal” and can identify deviations from those norms. More significant deviations could imply impact to a customer.

For count valued features, we base our approach on common natural language processing (NLP) techniques. The counts of the more common errors/logs over a time window can indicate the intensity of a problem, and the existence of rare keys can indicate a complete service impairment. Thus, common NLP techniques are applicable to our use case and provide a framework for creating features (embeddings) from which a model may learn. Each error code or system message indicator is treated as a “word” and

passes this context to our feature pipeline as a “bag of words”. The data is then transformed via a pre-trained TF-IDF pipeline before using it in the model.

Our feature processing pipelines use classes from sci-kit learn. For example, all features are normalized before training using preprocessing classes. Continuous features are standardized and (sparse) count features are scaled by their maximum absolute value. In addition, the TF-IDF embeddings are implemented using the corresponding classes found in scikit-learn for this purpose.

9. Results

By utilizing the machine learning model, we achieved a 40% improvement in customer engagement, as measured by the abandon rate of sessions (when a customer opens the application and leaves without clicking or asking a question). This was measured through an A/B test where all customers selected to get a card were split into two groups: the treatment group who were shown the card and a control group where we held back the card. This business impact was accompanied by a measurable increase (~42%) in the click-through rate of cards when compared to simple card display rules based on telemetry cutoffs. This improvement was calculated by comparing each approach to a baseline over their respective time periods in production to control for temporal differences. The final number is the relative difference between the two statistics after controlling for the baseline.

Using the general architecture depicted in Figure 8, with simple horizontal scaling, the platform can process very large amounts of incoming telemetry data. When converted to feature sets, the production systems are handling 100k transformations per second. This is using feature windows that vary between 6 and 24 hours based on the dataset and requirements for the model. This scaling is nearly seamless if the code is written with good microservice design principles in mind, and the ease of increasing replicas for an application is handled. Kubernetes also supports elastic scaling based on certain metrics, which is useful for processing spikes or bursts in traffic.

Scale aside, researchers and ML engineers have the ability to deploy models and feature pipelines as needed. This reduces code rewrite from POC or research code to production code as the production systems support many languages, ML model types and complex pipelines. Additionally, troubleshooting and operations (MLOps) is greatly aided by Prometheus integration with thresholding from Grafana and other tools to ensure problematic components are dealt with quickly. System availability is high due to the redundant nature of Kubernetes and requiring all components have a minimum of two replicas and dual components when they lie in the critical path.

10. Looking Ahead

These sorts of open source platforms and architectures are simply the beginning of a sweeping change coming to not only telecommunications, but every other industry. Big telcos certainly have an advantage now, given the vast datasets being collected now or in the future. While the physical properties of components and systems are well understood, there is always room to identify new correlations in failure or impairment types. With new hardware, more optical components and higher speeds -- and the fact that connectivity is so vital to everyone -- this shift to self-identifying and self-healing networks will become more critical. While ML models are not specifically required for every problem, the collection and analysis platforms are, because the data is required to work towards more data-driven decision making. Also, in the future, the platform output itself will be fed through various anomaly detection model types to determine if system performance has deviated sufficiently to warrant human interaction. Platform failures and other behaviors can be modeled to increase self-healing, predictive scaling (as opposed to reactive) and automated root-cause analysis -- down to the specific component.

These platforms will need to adapt in the coming years, to leverage FPGA or GPU hardware to handle increasingly more complex analysis and ML workflows. Many companies are already seeing the benefits from GPU, though TPUs and other hardware will be available to further drive down these large costs.

11. Conclusion

Again, the core problem here is that customer service and many other applications may require ML to make a measurable impact in customer experience. As we have shown, not only were many open source technologies instrumental in building the solution, they were not the only important pieces. These systems may need to scale which will require deep dives into some of the core components that get executed many times over, and those should be optimized as much as possible. While choosing Python (the same language as is common for research) initially proved troublesome, performance has increased dramatically and new avenues to optimize these functions, such as GPU offload have come a long way.

When thinking about the model, effort is required to determine if ML is necessary at all and how it can make a positive impact on the problem. For this problem, a contextual bandits approach was chosen but that does not mean is it the correct solution for other similar problems.

By now, after reviewing the problem, solution and challenges to solving this problem at large scale, it is hoped that this information is not only helpful but the start of a change in thinking about how problems will be solved going forward into the next decade. ML is not beginning but continuing to alter many parts of our organizations and this will only accelerate. The information presented here may be somewhat specific to telco telemetry data, but in fact it can be applied to most any similar dataset.

Our customers are the core of what we do and drive us to strive to do better in terms of not only providing services but help when those services become impaired. This solution and others not only aids in our graceful and accurate handling of these, but the hope is for all of us to continue to think about these problems from the customers perspective.

Abbreviations

CM	cable modem
CMTS	cable modem termination system
CPE	customer premise equipment
DNN	deep neural network
FPGA	field programmable gate array
GPU	graphics processing unit
HSD	high speed data (internet service)
IoT	internet of Things
JSON	JavaScript object notation
OS	operating system
MAC	media access control address
ML	machine learning
MLOps	machine learning operations

NLP	natural language processing
POC	proof of concept
POP	point of presence
RDK	Reference Design Kit
RSSI	received signal strength indicator
STB	set-top box
TF-IDF	term frequency, inverse document frequency
TPU	tensor processing unit
VBS	video backend services

Bibliography & References

- [1] *Data Over Cable Service Interface Specification Proactive Network Maintenance*, PNM Best Practices Primer: HFC Networks (DOCSIS® 3.1) CM-GL-PNM-3.1-V01-200506, 05/06/2020, <https://www.cablelabs.com/specifications/CM-GL-PNM-3.1>
- [2] *What is Customer Experience*, Blake Morgan; Apr 20, 2017, <https://www.forbes.com/sites/blakemorgan/2017/04/20/what-is-customer-experience-2/#8b5958170c2b>
- [3] Reference Design Kit (RDK), <https://rdkcentral.com>
- [4] *Protocol Buffers*, <https://developers.google.com/protocol-buffers>
- [5] *Apache Arrow*, <https://arrow.apache.org>
- [6] *Apache Kafka*, <https://kafka.apache.org>
- [7] *Docker, Inc.* <https://www.docker.com>
- [8] Kubernetes, Linux Foundation, <https://kubernetes.io>
- [9] Peng Kang, Finding Kafka's throughput limit in Dropbox infrastructure, Jan 30, 2019, <https://dropbox.tech/infrastructure/finding-kafkas-throughput-limit-in-dropbox-infrastructure#:~:text=This%20result%20provides%20guidelines%20for,throughput%20of%20future%20use%20cases.>
- [10] *Apache Flink*, <https://flink.apache.org>
- [11] *Apache Beam*, <https://beam.apache.org>
- [12] *Seldon Core*, Seldon Technologies Ltd, <https://docs.seldon.io/projects/seldon-core/en/v1.1.0/>
- [13] *Feature Scaling*, https://en.wikipedia.org/wiki/Feature_scaling
- [14] CITE-WIFI *Observing home wireless experience through wifi aps*. Ashish Patro, Srinivas Govindan, and Suman Banerjee; In Proceedings of the 19th Annual International Conference on Mobile Computing & Networking, MobiCom '13, page 339–350, New York, NY, USA, 2013. Association for Computing Machinery.
- [15] CITEA *The nonstochastic multiarmed bandit problem*. Auer, Peter, Bianchi, Nicol'o C., Freund, Yoav, and Schapire, Robert E.; SIAM Journal on Computing, 32(1):48–77, 2002.
- [16] CITEB *The epoch-greedy algorithm for multi-armed bandits with side information*, Langford, John and Zhang, Tong; In Advances in Neural Information Processing Systems 20, pp. 817–824, 2008.
- [17] CITEC *Playing Atari with Deep Reinforcement Learning*, Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller; <https://arxiv.org/abs/1312.5602>

- [18] CITED *Learning from Logged Implicit Exploration Data*, Alex Strehl, John Langford, Sham Kakade; 14 Jun 2010, <https://arxiv.org/pdf/1003.0120.pdf>
- [19] CITEE *Batch Learning from Logged Bandit Feedback through Counterfactual Risk Minimization*, Adith Swaminathan, Thorsten Joachims; 2015
https://www.cs.cornell.edu/people/tj/publications/swaminathan_joachims_15c.pdf
- [20] CITEF *A Contextual-Bandit Approach to Personalized News Article Recommendation*, Lihong Li, Wei Chu, John Langford, Robert E. Schapire; 1 Mar 2012, <https://arxiv.org/pdf/1003.0146.pdf>
- [21] CITEG *Top-K Off-Policy Correction for a REINFORCE Recommender System*, Minmin Chen, Alex Beutel, Paul Covington, Sagar Jain, Francois Belletti, Ed H. Chi., 6 Dec 2018
<https://arxiv.org/pdf/1812.02353.pdf>
- [22] CITEH *Tutorial on Counterfactual Evaluation and Learning for Search, Recommendation and Ad Placement*, Thorsten Joachims, Adith Swaminathan; SIGIR 2016, 17.07.2016,
<http://www.cs.cornell.edu/~adith/CfactSIGIR2016/>

Thank You FCC -- Now We Have 1.2 Ghz of 6 GHz Spectrum so How Does the Cable Operator Utilize It?

A Technical Paper prepared for SCTE•ISBE by

J.R. Flesch

Director, Advanced Technology, HN/CPE
Commscope
3871 Lakefield Drive
jr.flesch@commscope.com

Bryan Pavlich

Staff SW Engineer
Commscope
3871 Lakefield Drive
678 473 8346
Bryan.pavlich@commscope.com

Kurt Lumbatis

Distinguished SW Engineer
Commscope
3871 Lakefield Drive
678 473 2921
Kurt.lumbatis@commscope.com

Charles Cheevers

CTO/CPE and HN
Commscope
3871 Lakefield Drive
678 473 8507
Charles.cheevers@commscope.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Testing the 6E Proposition for Home Indoor Deployment	3
2.1. Wi-Fi Test House	3
2.2. Analytical Expectations	4
2.3. Measurements and Reconciliation	11
2.4. Implications for 6E Product Planning	12
3. Coming Attractions	17
3.1. 320 MHz Channel Bandwidth.....	17
3.2. 8 dBm/MHz PSD	18
3.3. 4096-QAM	20
4. 6E Air Time Considerations	21
4.1. AP Framing	21
4.2. The Single Radio Extender (Repeater) Dynamic.....	22
4.3. The Rectangular Mapping Challenge	22
4.4. Queues, Priority and Fairness.....	24
5. Conclusion.....	25
Abbreviations	25
Bibliography & References.....	26

List of Figures

Title	Page Number
Figure 1 - Commscope Wi-Fi house Floorplan	4
Figure 2 - UDP Bitrate Curves @ 5 dBm/MHz.....	5
Figure 3 - AP to Client UDP Bitrates.....	6
Figure 4 - Bitrate “heat map” for Wi-Fi house from AP in Pool Room	7
Figure 5 - Improvement in Whole House Path Losses with Midpoint AP	8
Figure 6 - Wi-Fi House Bitrate Heat Map from Midpoint AP	9
Figure 7 - Wi-Fi House Heat Map for 5 MHz, Wi-Fi 6 from Midpoint AP	10
Figure 8 - Impact of Contention at 5 GHz	11
Figure 9 - Measured Bitrates Throughout Wi-Fi House	12
Figure 10 - 6E Wireless Link for Gaming	13
Figure 11 - 6E Upgrade to Existing Dual-band AP at WAN.....	14
Figure 12 - Reconfigurable Multiband AP	15
Figure 13 - Mux-less Alternative for Midpoint AP.....	16
Figure 14 - 6 GHz Wi-Fi Spectrum (with 320 MHz channels).....	17
Figure 15 - Effect of Increasing PSD to 8 dBm/MHz	18
Figure 16 - Comparative Reach to Clients from Mid-Home AP @ 5, 8 dBm/MHz.....	19
Figure 17 - Measured Bitrate Improvement at 8 dBm/MHz (vs 5)	20
Figure 18 - Mu-PPDU Behavior at Scheduling Onset.....	22
Figure 19 - Fixed Frame Mapping via Progressive Largest Area	23
Figure 20 - Home 6E Network Proportionally Fair Mapping of Mu Data	24

1. Introduction

From a Wi-Fi perspective, Christmas arrived in March of this year with the FCC greenlighting LPI exploit of the entirety of the 6 GHz band, to a fixed PSD of 5 dBm/MHz – and this, without the complication of implementing an AFC (Automated Frequency Controller) system to dynamically allocate EIRP vs channel masks for devices. A perhaps unappreciated corollary is that only the Wi-Fi 6 MAC and beyond are permitted to be utilized by devices contending for the 1.2 GHz worth of spectrum, which will promote much more deterministic and efficient wireless delivery metrics. AFC -- and access to additional EIRP footprint --, have been kicked down the road for the time being; in the meantime, SPs have the opportunity to pursue a glut of new bandwidth opportunity with a still-significant indoor radiation footprint.

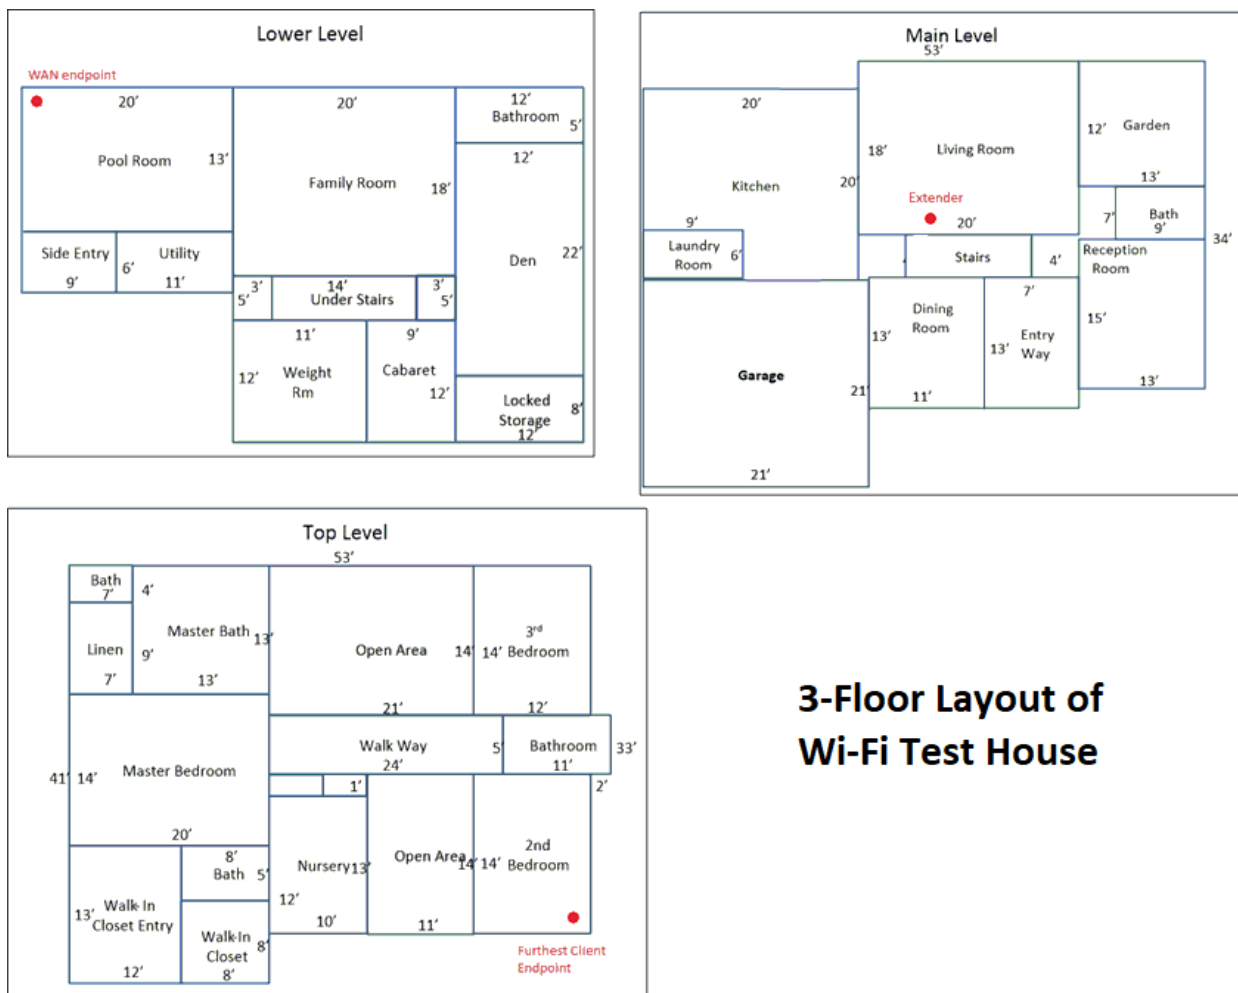
Preliminary coverage performance of 6E (Wi-Fi 6 @ 6 GHz) reference platforms will be presented in this paper, along with recommendations on initial leverage of the spectrum which should produce better bitrate and latency coverage of a dwelling – along with plug-and-play upgrades to extant home networks which minimize user installation headaches. There is a brief section outlining opportunities for early creep-in of features formally scheduled for Wi-Fi 7 and some additional test data (unfortunately limited by transmit chain noise fidelity) hinting at the benefits of raising the PSD to 8 dBm/MHz. Finally, there is a closing section which anticipates the adoption curve for various aspects of the scheduling functions which are due to sit just above the Wi-Fi 6 MAC, along with an analysis of some of the tradeoffs which are beginning to emerge as the radio, FEM and filter vendors balance time-to-market against both MAC accessibility for tuning performance and analog circuit challenges on the RF chain front.

2. Testing the 6E Proposition for Home Indoor Deployment

Two particular bitrate reach propositions begged initial testing: a) how far can you space 4x4 trunk endpoints in a “typical” home setting given 5 dBm/MHz PSD and 160 MHz of channel BW at 6 GHz?; and b) given the likely configuration of clients as 2x2 devices, what is the equivalent “throw” for bitrate coverage between a 4x4 AP and 2x2 client in such a setting? For that matter, how do these “service reach” numbers compare to 5 GHz performance (at higher PSD, lower BW and with airtime contention losses due mainly to heterogeneous MAC accesses to the same medium)?

2.1. Wi-Fi Test House

We were fortunate to have available a 3-level, 5000 sq ft class Wi-Fi test house and pre-production reference platforms (capable of Wi-Fi 6 operation in both the 5 and 6 GHz bands) from two vendors to provide test telemetry which we could compare to analytical expectations. The three floors of the test house floorplan are shown below:



3-Floor Layout of Wi-Fi Test House

Figure 1 - Commscope Wi-Fi house Floorplan

Note the test endpoint placements, chosen to produce challenging pathlosses which could provide inference for more modest layouts as well (the average US home being ~ 2700 sq ft; average apartment in the vicinity of 1000 sq ft). Expectations were that proper extension of a 4x4 6E trunk from WAN insertion at the lower level pool room to approximate midpoint of the test house would allow us to paint 2x2 client TCP bitrate coverage everywhere in the home to between 1 and 2 Gbps (presuming leverage of 160 MHz channel BW and a PSD of 5 dBm/Mhz).

2.2. Analytical Expectations

The basis for the confidence in coverage stems from raw UDP throughput analysis. The case of 4 spatial stream leverage of varying channel BW's for a symmetric 6E link (versus accumulating path losses) is shown below:




Figure 2 - UDP Bitrate Curves @ 5 dBm/MHz

Note that the abscissa lays out losses past 1 meter off of the antennae. It is also noteworthy that the general shape of the MCS breaks (bitrate steps) is exactly mirrored across all BW – this is due to EIRP being metered exactly as a function of occupied BW (fixed PSD). In terms of actual link performance, these UDP asymptotes will be reduced by TCP overheads (typically 10-20%, depending on packet parameters) and will also be subject to implementation losses (systemic, as analog RF chain fidelity issues – principally noise and filter shape; and circumstantial, as spatial stream compromises – polarization and element diversity losses).

With respect to asymmetrical 6E link performance, a similar UDP bitrate waterfall can be estimated for the case of a 4x4 AP linking with a 2x2 client:




Figure 3 - AP to Client UDP Bitrates

Note the callout on pathloss @ 62 dB: this is the estimated LOS transit loss from the Wi-Fi house pool room up to the furthest potential client location in the upper floor 2nd bedroom. The budgetary accounting includes FSPL and multiple drywall and floor transitions (the former measured at 4.5 dB per wall and the latter, 9 dB per floor – in the 6 GHz band).

Using this single-vector predictor for pathloss, the following heatmap for bitrate support (as 4x4 AP linking to 2x2 client at 160 MHz BW, 5dBm/MHz PSD) by room is produced:



Figure 4 - Bitrate “heat map” for Wi-Fi house from AP in Pool Room

The three floorplans stack one above the other, so as anecdotal experience indicates, as you move along a 3D diagonal from the WAN insertion point in the pool room lower level and progress towards the far 2nd bedroom on the top level, the predicted bitrate support “cools”. The analysis suggests that the challenge of supporting > 1 Gbps in a large home from a single 6E GW device located at the WAN connection point (typically against one wall in a lower level) would not be possible. In particular, user endpoint uses in upper floor bedrooms appear underserved.

This layout challenge begs for link extension. If a spatial midpoint were chosen in the home, one would expect to be able to paint better coverage at the home extremes via reduction in the worst-case pathlosses posed. Graphically, the placement of a 6E extender at home midpoint (mid-level living room chosen for the Wi-Fi house) produces the following improved link parameters:




Figure 5 - Improvement in Whole House Path Losses with Midpoint AP

The modeling suggests we can have our > 1 Gbps large home by 6E extension at home midpoint with leverage of the 5 dBm/MHz PSD and 160 MHz channel BW. Here's what the modified heat map produces:



Figure 6 - Wi-Fi House Bitrate Heat Map from Midpoint AP

If the extender is multiband (a sensible enough proposition, given the foregoing conclusion of legacy clients) then we can also predict coverage at 5 GHz, using a Wi-Fi 6 MAC on a 4x4 AP capable of 1W EIRP at 80 MHz BW (but restricted to 65% airtime access due to contention):




Figure 8 - Impact of Contention at 5 GHz

These results beg some questions on extender band profile, spatial stream assignment and RF chain management (including antennae multiplexing, perhaps) for proper future-proofing of client band and MAC migration; we will tend to that in a later section.

2.3. Measurements and Reconciliation

Multiple days of testing were conducted at the Wi-Fi house to vet performances of the alternate vendor candidates and tune device alignments for repeatable data-taking in the various 6E use cases. The summary results for these follows:




Figure 9 - Measured Bitrates Throughout Wi-Fi House

The key takeaways here are that: 1) a NxN <-> NxN bookended trunk link from WAN insertion point in the basement to midpoint of the home could sustain an ~ 1.5 Gbps TCP bitrate; 2) worst-case link throughput for NxN <-> NxN just cleared 400 Mbps for the long diagonal reach from WAN to 2nd bedroom client (MUCH better than the single-vector analysis predicted); and 3) 2x2 client link service from an NxN AP in the home midpoint (living room, mid-floor) was sustainable everywhere to ~ 1 Gbps. A secondary observation is that, in the client link data, the implementation loss looks quite small (< 10% off of target expectations once UDP has been discounted for TCP overhead) – but in the full trunk bookend case, where one might expect double the spatial stream efficiency to show, an average multiplier of 1.44 (instead of 2) was obtained (~ 25% implementation loss). This points to the general observation that full spatial diversity is easier to obtain in the oversampled case than when one relies on full orthogonal recovery from all antennae in an N:N setting. (Losses due to antenna element blanking, insufficient packaging diversity and polarization mismatches are harder to mask).

There was an additional lurking issue which only became more pronounced when we attempted to vet the benefit of a PSD boost from 5 dBm/MHz up to 8 dBm/MHz. Additional detail on this is captured in subsequent sections.

2.4. Implications for 6E Product Planning

Two new product opportunities seem to immediately exist in bringing 6E CPE to market which do not rely on simultaneous emergence of clients (the “chicken and egg” marketing conundrum when new band exploits appear). The first such opportunity would be a tri-band Wi-Fi 6 AP (2.4/5/6 GHz) as a “ready for Wi-Fi 6” play for consumers looking to upgrade or create a home wireless network. This anticipation play has as its rationale the ability to mount a Wi-Fi 6 MAC in all bands and so both continue to serve legacy clients (in whatever native MAC they utilize) and provide improved connectivity for both new multiband Wi-Fi 6 clients and specific new 6E devices as these emerge in the market. The second device

would appear to be a bridging device (Ethernet to 6E) which could be applied against several “touchless” home wireless network upgrades. Some use case examples follow.

A key marketing opportunity lies in the superior latency and contention-free performance of 6E. Not mentioned in prior sections but potentially holding the key for particular service mounts at 6 GHz is the notion that, until the rather prodigious link capacities of 160 MHz channel operation in the band are reached (and scheduling onset occurs in the AP), link latency easily outperforms that encountered in either of the 5 GHz bands or 2.4 GHz (as in: a sub-millisecond asymptote – essentially the time required to rasterize the arriving data against all available RU’s and dispatch it). Viewed this way, 6 GHz may be thought of as virtual wireline (which need not be pulled, obviously).

Applications particularly sensitive to latency include networked twitch games – the bulk of which are presently connected in the home wirelessly (by a factor of roughly 7:3 versus Ethernet). Moving these gaming systems (and PCs) to the low-latency 6 GHz band can be done by simply upgrading a GW AP with a 6 GHz radio (Tethered Ethernet dongle as the migratory – as in “no new AP” -- option) and doing the same for the gaming system or PC. Service reach to several hundred Mbps for even just two spatial streams seems guaranteed by our testing, regardless of where in the home lies the gaming device; and there would be no intervening alternate band/MAC backhaul to impress nondeterministic latency on the otherwise bridged link. Yes, not the cheapest solution – but as pointed out previously, no pulling of Ethernet either – and simple plug-in, “touchless” remediation of a perceived problem is always valid market rationale.

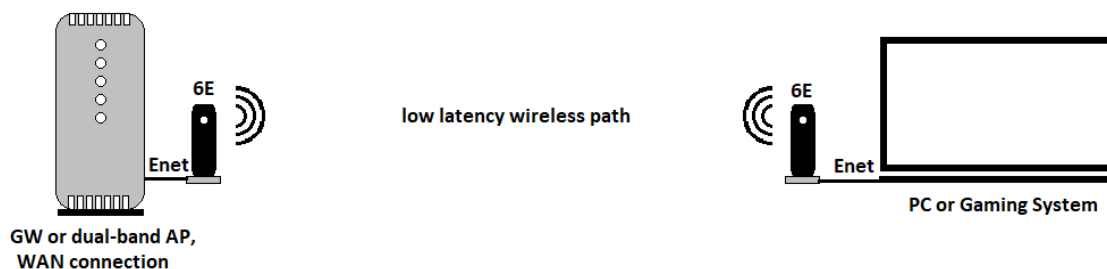


Figure 10 - 6E Wireless Link for Gaming

Another bookend 6E solution which need not wait on the emergence of explicit 6 GHz clients is the trunk backbone alluded to in the document’s prior sections (for those cases when floorplan service footprint overwhelms a location-marginalized single GW). Using 6E to extend a virtual WAN connection to the middle of the home – on first deployment being only responsible for trunk traffic haul between endpoints and leaving the entirety of client support to the 2.4 and 5 GHz bands – would greatly improve whole home coverage and eliminate Wi-Fi dead zones by collapsing the 2.4/5 GHz service radii to all clients (thus permitting higher sustained MCS and throughput for each of the links). As the client palette shifts over time, the load balancing of the trunk could move in concert – perhaps offloading unprioritized trunk traffic to 5 GHz as replacement clients migrate to the 6 GHz band. As with the gaming latency solution, the fix could involve a simple tethered 6E radio augmentation to existing GW and extender – or in the case of a single dual-band AP serving an entire dwelling, adding what amounts to a tri-band extender (2.4/5/6 GHz) to the home midpoint and plugging in 6E capability for the original dual-band AP.




Figure 11 - 6E Upgrade to Existing Dual-band AP at WAN

This segues into the notion of adaptable AP resource allocation – specifically with respect to analog circuit elements and antenna farm exploits. The multiplexing options lay out nicely if one allows for broadband-matched antennae and potentially, FEMs. Different price points will determine base multiband capability, but a specific example might drive home the point. Presuming that 2.4/5 GHz FEMs can be stretched to both 2.4/5 and 2.4/6 GHz devices (the built-in diplexers bifurcating the 2.4 from the other, common, band support), we already have the BAW/SAW devices which delineate the 2.4, 5 and 6 GHz bands (so-called “coexistence filters”) and so can consider our switching options. Hypothesizing an adaptable 8-chain AP leads to this, among other, possible implementations:




Figure 12 - Reconfigurable Multiband AP

This diagram describes the potential of multiband FEMs and configurable 5/6 GHz radio chips to implement an AP with reconfigurable band coverage – as might be advantageous when addressing the desire to first upgrade MAC capability in the lower bands before switching on actual 6E support for new client devices (as they come available). And while not evident in radio vendor planning at this stage, it is not inconceivable that configurable Wi-Fi 6 radios for 2.4/5/6 GHz appear some day (along with supporting infrastructure like wideband FEMs and triplexers). Clearly, every home’s traffic profiles will be at least slightly different, so to the extent that band and spatial stream flexibility can be designed in, such adaptability improves market edge.

A quick sidebar is in order. As a more philosophical discussion, the exploit of the 2.4 GHz ISM band for Wi-Fi might be ripe for reconsideration given the advent of the many, unfortunately disparate, IoT radio mesh exploits which target that same spectrum (Zigbee, Thread, BLE). Coordinated exploit of the band by the many interested parties is a bit brute force and clumsy at this stage and perhaps a graceful exit (or at minimum, an explicit band-sharing FDM solution) may be the future play for 802.11. The flip side of the coin may also show forth: IoT applications normally relegated to other radio technologies may find new homes in Wi-Fi 6 narrowband assignment. To the extent that constrained end devices could host a Wi-Fi 6 stack, this bears watching.

It is important to bookmark one additional item in regards to this: the potential for early adoption of 320 MHz channels in the 6 GHz band (3 such ought to be available in North America). Assuming RF chain fidelity can match the offering, the PHY bitrate asymptote for a 2x2, 320 MHz link would be 4.8 Gbps. Adoption of this very broadband channel capability might provide sufficient rationale (and capacity insurance) to park initial 6E AP capability at 2 RF chains. This greatly simplifies the multiplexing overhead in our hypothetical early adoption 6E AP, since the 6 GHz chain connects can be fixed at 2. If we also presume that 2.4 GHz AP coverage for the whole home is adequately served by our WAN entry point location, a lower cost alternative for Wi-Fi 6 extension might take on the following form:

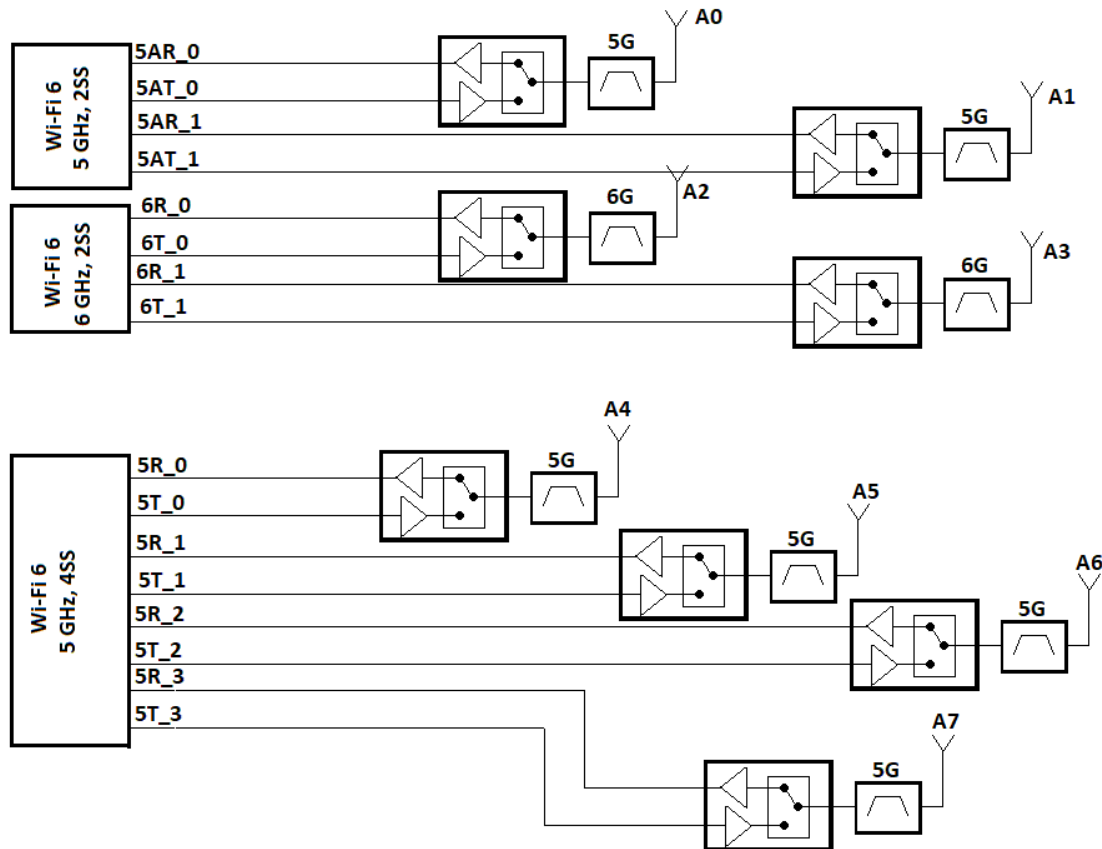


Figure 13 - Mux-less Alternative for Midpoint AP

Note that this type of extender commits itself only to the upper bands (assuming whole-home 2.4 GHz coverage from the WAN insertion point would be adequate) and avoids the multiplexing costs of our prior example. At issue is the degree of future-proofing offered by a 2-chain 6E (potentially 320 MHz BW) capability; if deemed adequate, this makes for a less complex and cheaper solution. The premise is that you have one, four-chain “heavy lift” capability at your choice of 5 GHz bands (for trunk hauling, say) and one lightweight assistant for proximate distance clients in that same band. The 6E capability, though spatially constrained, has a very flexible application capability (especially given the initial paucity of 6E clients and the prospect of neatly commanding all significant data exchanges from the AP itself).

At this point, it seems incumbent to also point out that the broad swath of uncontended 6 GHz BW avails itself of a notion which outlived its usefulness in the lower bands due to issues with airtime oversubscription and heterogeneous MAC accesses: the repeater. Simply put, one could mount a 6 GHz

repeater as an AP mid-home and orchestrate air time on that single 6E radio for what would function as front- and backhaul between otherwise farflung and bitrate-lean clients. Such an arrangement could be implemented in straightforward fashion with an AC-plugged “wall wart” and might represent the means by which either distant clusters of 6E clients – or a bit-hungry single 6E service mount in a detached shed or garage – gets its high bitrate and low latency connection to the home edge WAN located an appreciable distance away. Prior instantiations of the repeater function required explicit diplexing of the repeat trunk, essentially replicating air time losses on an additional channel. While merely inconvenient and problematic in the 5 GHz band, such implementation at the grossly oversubscribed 2.4 GHz band meant such attempts were all but doomed. 6E’s expansive spectrum would allow many of these private “repeat” meshes to co-exist without extracting a penalty on the “main” channel.

3. Coming Attractions

The ink on the March FCC announcement laying out the access to a new chunk of spectrum for Wi-Fi has not even dried and there are already three areas of bitrate service growth which are garnering attention. Two of these are related to scaling features planned for Wi-Fi 7 (320 MHz channel BW and 4096-QAM) – and the other has to do with unlicensed spectrum proponents wishing to migrate to a slightly higher PSD (8 dBm/MHz vs the present 5). All of these should yield opportunities for link improvement but there is also a clear benefit hierarchy to the proportional improvements one might expect. These are covered in decreasing magnitude of impact in the following sections.

3.1. 320 MHz Channel Bandwidth

By far the most attractive of the features-in-waiting is the ability to accommodate 320 MHz channel BW. The scale benefit here is obvious: double the maximum bandwidth and you will potentially yield double the supported bitrate. As alluded to in a prior section, this also gives us the option of considering fewer spatial leverages for the capability, since the raw capacity of even a single stream (2.4 Gbps – at admittedly very short service throws) dwarfs performance of 4x4 Wi-Fi systems in the lower bands.

But 320 MHz brings its own challenges as well: due to achievable filter Q’s and proximate band signaling, there are likely only 3 such channels available across all of the 6 GHz spectrum (and perhaps only 1 in Europe). And the very aperture of these large channels make them susceptible to rogue or OBSS interferers – which coupled to the dynamic range requirements for dense QAM constellation leverages, could make exploiting the capability a difficult proposition “in the wild” (especially with respect to forced concentrations of disparate SSIDs – as happens in MDUs. In short, it is a very worthy feature to have, but perhaps not fully leverageable in all AP deployments.




Figure 14 - 6 GHz Wi-Fi Spectrum (with 320 MHz channels)

3.2. 8 dBm/MHz PSD

We have already shown fairly decent whole home coverage at the FCC's original setting for LPI services (5 dBm/MHz). Pushing up the PSD does exactly what the bitrate waterfall curves suggest: you potentially can establish a link MCS higher by a step (or two) and reap the scaled bitrate benefit for doing so. The following 5 dBm/MHz chart (for 4x4 6E) indicates potential benefit (at 160 MHz BW) at one particular path loss operating point:




Figure 15 - Effect of Increasing PSD to 8 dBm/MHz

If we then take our original heat map and recalculate the room-by-room reach at 8 dBm/MHz, we get the following comparative results (5 dBm/MHz on the left and 8 dBm/MHz on the right):

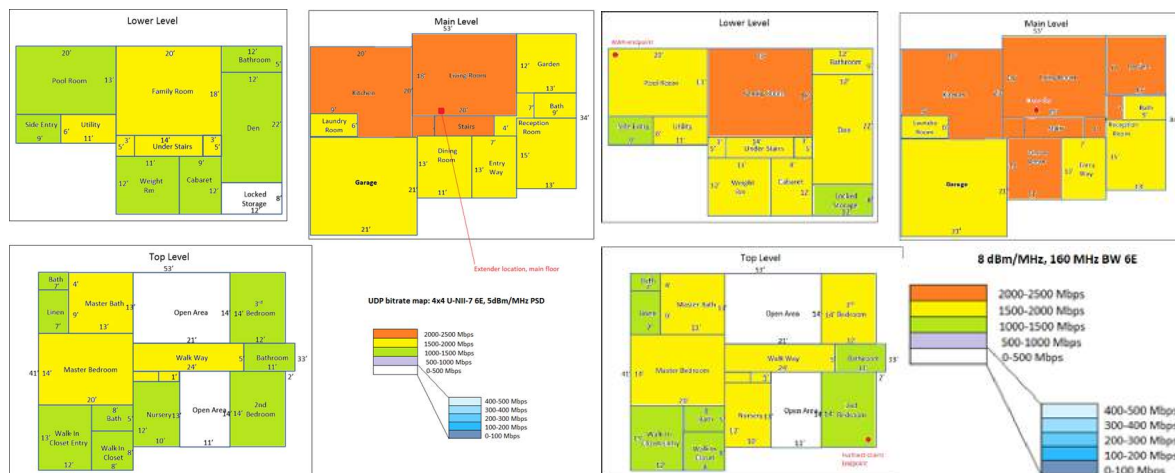


Figure 16 - Comparative Reach to Clients from Mid-Home AP @ 5, 8 dBm/MHz

Note the expanded reaches for the 2 Gbps+ and 1.5 Gbps+ regions for the higher PSD setting.

An attempt was made to test this analytical promise but the effort required upgaining antennae on a vendor-supplied reference platform (the standard 2 dBi elements and RF PA chain maxed out at an EIRP of 27 dBm, 3 shy of that necessary to validate 8 dBm/MHz at 160 MHz BW). Since we wanted to explore max impact (as opposed to doing a scaled test at lower BW), the decision was made to upgain the antennae elements to buy the necessary EIRP. The exercise produced mixed results, however. As it happened, receiver-commanded power back-offs at any path loss much less than full span of the Wi-Fi house meant the transmit side could not reach the desired EIRP. The fundamental culprit appeared to be PA fidelity (excessive noise); for the lone case where we could validate the bitrate improvement, the benefit for the 3 dB extra PSD showed as ~ 32%:

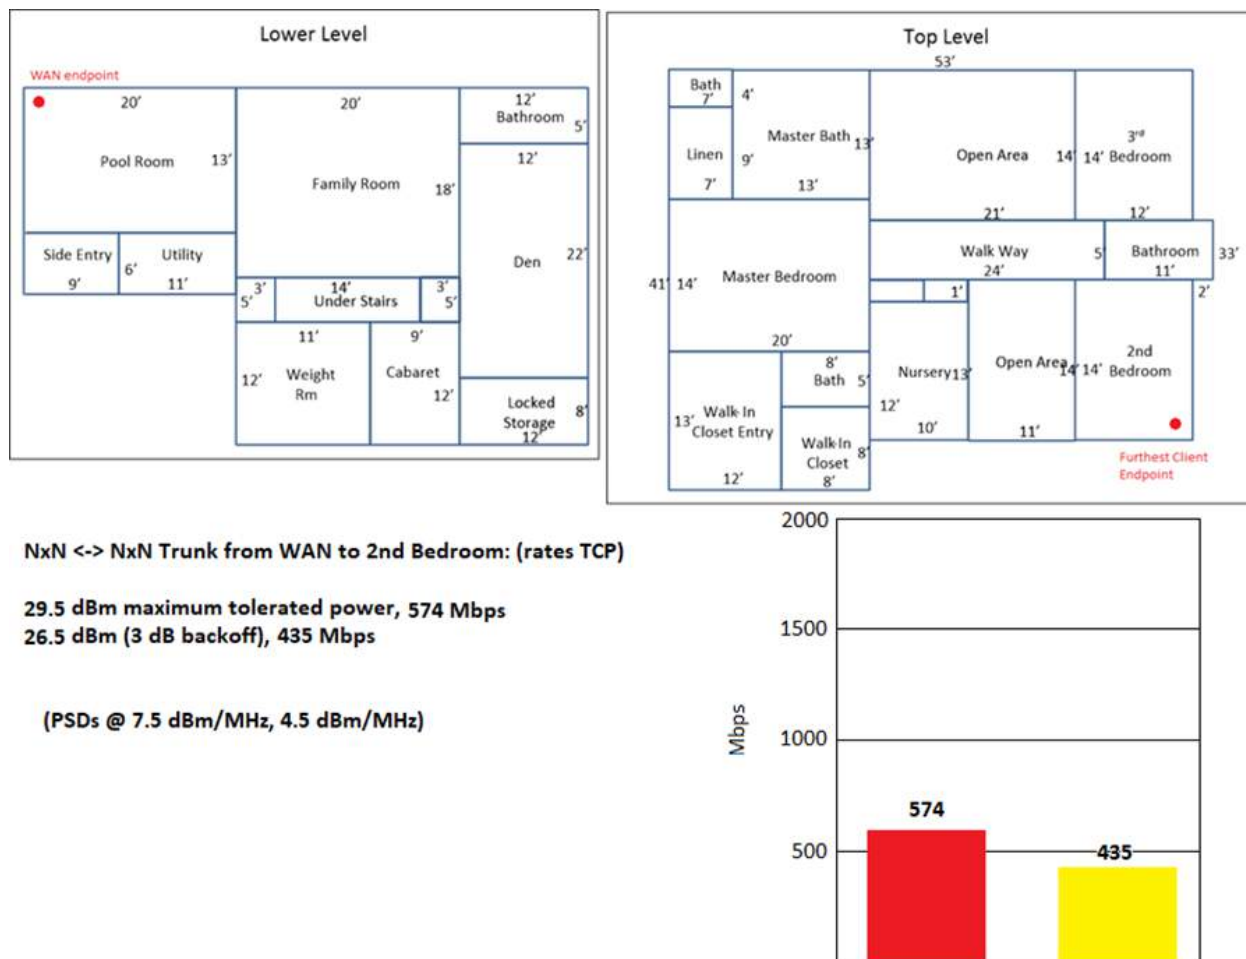


Figure 17 - Measured Bitrate Improvement at 8 dBm/MHz (vs 5)

3.3. 4096-QAM

Ranking last in terms of desirable upticks to signaling bitrate is 4096-QAM – not that such spectral efficiency is not laudable but that, relative to 1024-QAM, it places extreme implementation pressures on the analog circuit goodness of hosting implementations for perhaps not a great operational reward (in terms of signal throw). The fundamental benefit boils down to a 20% improved spectral efficiency given the raw symbol length of 12 bits versus 1024's 10. You might react with “Hey, gain is gain – so why not?”

But the rub is that, to deliver such a dense constellation, the link needs to stand up an SNR which approaches 40 dB. If you attempt to do that for a 160 MHz channel, the napkin math is striking: only signals above ~ -50 dBm need apply. This pretty much reserves the use of 4096-QAM to same-room-only (relative to the AP) and multi-spatial stream implementations -- and even at that, clock phase noise and transmitter chain fidelity need to be pristine (so as not to pre-emptively rob SNR and reduce the already constrained service radius). The spectrum efficiency is remarkable (9 bps/Hz for ¾ rate coding) – but the exploit aperture makes this feature a bit of a unicorn.

4. 6E Air Time Considerations


The Wi-Fi 6 MAC contains five particularly noteworthy framing aspects (Trigger Frames, Block Acks, Buffer Status polling, TWT and MU-OFDMA) which streamline wireless community accesses to the medium. Applied to a wide channel BW, these AP-managed on-air exchanges can mimic data arrival heuristics to an impressive bitrate threshold before queueing techniques (as for differential service data priorities) need be invoked and introduce something other than simple dispatch latencies to the delivery time for the packets involved. Air time access in general will move between variable-width data exchange windows (scaled via the amount of data to be moved in order to empty buffers) and framed exchanges whose duration defines a maximum performance latency target for prioritized traffic. In all cases, there is algorithmic art performed which seeks to exploit the rectangular capacity posed by RU's x frame time. And all this need be exhausted before resort of a second AP radio (channel) to overlay an FDM option to the proceedings should be considered to alleviate single channel exhaustion.

The general scheduling opportunity in Wi-Fi 6 decomposes along two fronts: the scheme for mapping RU's in the OFDMA matrix to fit the AP's client demands and the fairness doctrine applied to whatever queueing scheme is adopted to stage data for dispatch.

4.1. AP Framing

A high-altitude summary of the framing structure for 6E communications – which is commanded exclusively by the AP, is in order. Essentially, Trigger Frames are used by the AP to coordinate single user, multi-user and upstream PDUs on the channel in question. The ability to coordinate multi-user exchanges via OFDMA mapping is also critical to the higher efficiencies in Wi-Fi 6, since simultaneous access to the channel can be synchronized across all clients. (Mu-MIMO offers another exploitable axis – spatial diversity -- but though this is specified, Wi-Fi MAC implementations to date have eschewed this as overwrought).

In addition to the data transfers, status details and general housekeeping can be coordinated through triggered responses for the multi-user environment as well. Simple Block Acks to close out TCP requirements fall into this category. In addition, by soliciting Buffer Status Reports from clients, the AP has access to a key metric for determining scheduled queue servicing priorities based upon its knowledge of its own queues, all the client queues and target scavenging bitrates represented by the applied MCS @ RU profile assigned to each particular client. With this knowledge, any of the proportionately fair (PF) scheduling algorithms (M-LWDF or EXP-PF -- common to LTE and WiMAX OFDMA networks) may be employed to serve clients should brute-force FIFO rasterization of the outbound (or inbound) data start to produce framed exchanges that approach the maximum targeted frame size for the given channel bandwidth (scheduler onset). A reasonable bound for this might be something on the order of 5 msec or so.



General Schema for Mu-Data growth to drive Scheduling Onset

Figure 18 - Mu-PPDU Behavior at Scheduling Onset

4.2. The Single Radio Extender (Repeater) Dynamic

In one of our analytical cases, we describe the deployment of an extender AP to provide a mid-home anchor point to virtually extend the home WAN attachment to a more convenient radiating point for balanced support of 6E clients. This scenario does introduce an air time complexity for the trunk link, summarized by the realization that the single channel capacity calculation necessarily involves both the 4x4 hauling trunk and multiple 2x2 client links. Both of these capacities impinge on the single channel air time, with clients attached to the extender essentially consuming two blocks of air time (to/from the midpoint, with repeated to/from the WAN endpoint for backhaul/fronthaul) and with clients associated with the WAN side only accounting for single transfers each. The air time hit is not necessarily double for the extender clients, however, since the trunk capacity is defined by 4 spatial streams @ the MCS for the trunk and the client impact is essentially 2 spatial streams at a different extender/client link MCS (might be higher or lower than the trunk MCS and depends on client distance from the extender). The key is that the higher the sustained MCS to the clients can be made, the less demand on airtime to service them; and the overriding goal is air time efficiency (clients / frame).

4.3. The Rectangular Mapping Challenge

Because LTE and WiMAX have already tilted these windmills, there is a glut of prior art describing different schemes for assigning RU's to client service needs in very demanding capacity scenarios (typical client counts going into the hundreds). (Refer to the reference section.) Such is not likely to materialize for 6E home networks for the foreseeable future and there is value in simplifying some of the more elaborate schemes from those mobile networks to suit the more limited scale of home meshes (to cost benefit, or marginalization of firmware complexity). To put it more bluntly, SU data packaging – even if concatenated for several (few) users, looks to be the dominant initial scheduling tactic for 6E's emergence. This involves simply rasterizing the outbound data over full exploit of the channel BW available (as RU's) until buffer exhaustion – and repeating the behavior for other clients when such data intermittently arrives. Multiple simultaneous browser sessions describe this type of scenario.

However, once streaming applications get mounted (especially in tandem), MU data dispatch dominates and actual stepped rectangular fitment begins to be necessary. In broad strokes, data bound for dispatch to multiple clients can be represented as a strip of symbols – such symbols the result of the achievable MCS on its own link. This strip, taken as a rectangle of area 1, can be applied in raster fashion across

whatever span of RU's is considered fair. In those other networks, there can be many clients, each with different rectangular blocks to move – so the technique is to triage the blocks, from largest to smallest, and fit them in a fixed size frame (symbol times x RU's) so that a minimal number of gapped or unused RU's occurs. (The problem is NP-complete, so a perfect or even optimum footprint is never guaranteed, but comparative performances can at least be evaluated for “consistent fit excellence”). Note that the problem is usually bound by a maximum RU assignment limit per client.



Figure 19 - Fixed Frame Mapping via Progressive Largest Area

In home 6E networks, the lower overlaid client count makes this more of a factorization problem. If you presume you are going to burst to perhaps 4 simultaneous clients, each of a different block size, you will likely want to rank-order the blocks with the area as a weighting factor. The sum of these weighting factors should map into the RU-span (BW) of the channel in question. By taking the area weights to assign the available RU's, you essentially get the 4 rasterized areas to be roughly the same length – which minimizes lost transmit opportunity and packs the burst as tightly as can be done. (Note that RU minimum grouping size will distort this somewhat.)

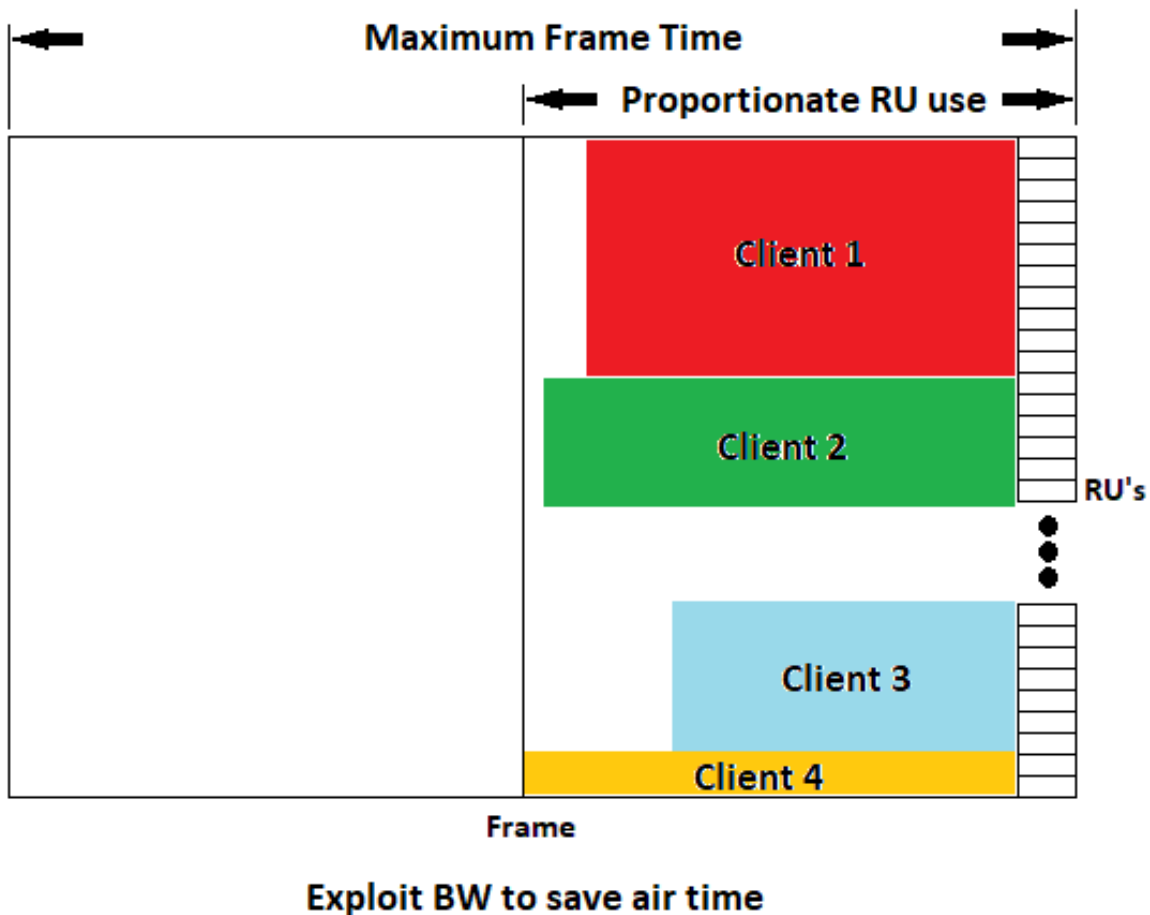


Figure 20 - Home 6E Network Proportionally Fair Mapping of Mu Data

4.4. Queues, Priority and Fairness

It is becoming painfully evident that radio vendors will be reserving details of the rectangular mapping algorithm – along with frame sizing – for themselves to implement, leaving little tuning capability for users aside from QoS designation. The number of QoS levels will define the number of queues allowed (plus at least one shunt lane for pre-emptive priority – TWT -- or retransmissions). In LTE space, a typical algorithmic first resort is a move to two queues: “real-time” and “non-real-time” to be able to express a priority for traffic where latency and jitter are less tolerated (typically media streams). This simple bifurcation usually sees a reservation for somewhere between half and three-quarters of the link capacity for real-time consideration, with the balance (and the unused reserved) available for non-real-time traffic. This is likely to find its way into 6E networks – and stacks on top of the rectangular filling exercise above in the following way: instead of just one queue sourcing the blocks of data to be transferred, there are two or more (all but the last with a reserved area). The fitment problem is simply stepwise executed on successive priorities. First, all the top priority blocks get mapped against that reserved area, from largest area to smallest. If space is left over after priority traffic exhaustion, this is reclaimed for priority 2 (plus its own reserve). Priority 2 is mapped in similar fashion and whatever residual is left is used to accommodate the last priority. This exercise is truncated if the maximum framing area is reached before buffer exhaustion. In all cases, buffers are appropriately FIFO-shuffled (and the next frame is queued).

5. Conclusion

Wi-Fi's gifted new spectrum – with a MAC ante which sheds 20 years of multi-epoch, stacked and somewhat organic growth, is a shot of adrenaline for wireless indoor networks. The healthy LPI EIRP footprint of 5 dBm/MHz (without the complexity of seeking AFC database permissions and perhaps with more level yet to come) looks to provide Gbps+ coverage and low-millisecond signaling latency to even large floorplan American homes. Additional coverage insurance may come by way of additional channel BW in the not-too-distant future. With the Covid-19 pressures to WAN infrastructure apparently driving a new technology tranche there, the opportunity to craft even more demanding wireless home applications stand to be vastly improved. Certainly, media-infused experiences which immerse the user will be easier to mount – and with a much more robust signaling conduit to support IoT devices, expectations for next-generation remote applications – particularly security, telemedicine and aging-in- place, go up as well.

Eventually, there will be enough client demand that scheduled air time at 6 GHz will become a necessity – but this eventuality may still be several years away, given the multiple Gbps of capacity represented by diverse spatial exploit of the channel BWs involved. In the interim, lessons from OFDMA scheduling of LTE and WiMAX systems (802.16e) – which have been around for over 10 years -- may find themselves represented in 6E networks. The deferred onset of need for efficient scheduling algorithms means relatively boilerplate commutating exercises (weighted round robin) with perhaps only a single premium queueing structure (two total queues) might provide all the fair access consideration required for early generation 6E home network scheduling agents.

Abbreviations

AC	alternating current
AFC	automated frequency controller
AP	access point
BAW	bulk acoustic wave
BLE	Bluetooth Low Energy
bps	bits per second
BW	bandwidth
EIRP	effective isotropic radiated power
EXP-PF	exponential proportionately fair
FDM	frequency division multiplexing
FEM	front end module
FCC	Federal Communications Commission
FIFO	first in, first out
FSPL	free space path loss
Gbps	gigabit per second
GHz	gigahertz
GW	gateway
ISM	industrial, scientific and medical
LOS	line of sight
LPI	low power indoor
LTE	long term evolution

MAC	media access control
Mbps	megabit per second
MCS	modulation and coding scheme
MDU	multiple dwelling unit
MHz	megahertz
M-LWDF	maximum-longest weighted delay first
Msec	milliseconds
MU-MIMO	Multiple-user, multiple in, multiple out
MU-OFDMA	multiple-user OFDMA
OFDMA	orthogonal frequency division multiple access
PA	power amplifier
PSD	power spectral density
QAM	quadrature amplitude modulation
RF	radio frequency
RU	resource unit
SAW	surface acoustic wave
SCTE	Society of Cable Telecommunications Engineers
SPs	service providers
SSID	service set identifier
SU	single user
TCP	transmission control protocol
TWT	targeted wait time
UDP	user datagram protocol
W	watt
WAN	wide area network
WiMAX	worldwide interoperability and microwave access
6E	Wi-Fi 6 @ 6 GHz

Bibliography & References

eOCSA: An Algorithm for Burst Mapping with Strict QoS Requirements in IEEE 802.16e Mobile WiMAX Networks; Chakchai So-In, Raj Jain and Abdel-Karim Al Tamimi, Dept of Computer Science and Engineering, Washington University in St. Louis

IEEE P802.11ax/D4.2: *Draft Standard for Information technology – Telecommunications and information exchange between systems Local and metropolitan area networks – Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 1: Enhancements for High Efficiency WLAN*; 802.11 Working Group of the LAN/MAN Standards Committee of the IEEE Computer Society, Copyright 2019

A Novel Algorithm for Efficient Downlink Packet Scheduling for Multiple-Component-Carrier Cellular Systems; Yao-Liang Chung, Department of Communications, Navigation and Control Engineering, National Taiwan Ocean University, November 2016

OFDMA Downlink Burst Allocation Mechanism for IEEE 802.16e Networks; Juan I. del-Castillo, Francisco M. Delicado and Jose M. Villal’ on, Instituto de Investigaci’ on en Inform’ atica de Albacete, UCLM, Spain

The Power of Distributed Access Architectures (DAA)

Benefits of Digital Fiber Along with Remote-PHY

A Technical Paper prepared for SCTE•ISBE by

John J. Downey
Sr. CMTS Technical Leader
Cisco Systems
RTP, NC
(919)931-9453
jdowney@cisco.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Background	3
3. Deployment Scenarios	4
4. Remote PHY High Level Benefits	5
5. Digital Optics Benefits	5
6. Power and Space Savings	6
7. The Power of Digital and IP	7
8. Automation and Operational Support.....	8
9. “Real Life” Testing and Results.....	9
10. Future Gazing.....	11
11. Don’t Forget/Neglect the “Little Things”	12
12. Conclusion & Summary.....	13
Abbreviations	14

List of Figures

Title	Page Number
Figure 1 - Typical Analog Deployment.....	4
Figure 2 - Digital Fiber (DAA) Node Deployment.....	4
Figure 3 - DAA Shelf Deployment.....	5
Figure 4 - Typical Analog HFC Deployment	7
Figure 5 - Node Splitting in RF Domain	8
Figure 6 - Digital Combining/Splitting.....	8
Figure 7 - Lab Testing Diagram	9
Figure 8 - Real Network Distance Testing Diagram	9
Figure 9 - DAA for MDU Business Case.....	14

1. Introduction

The industry's latest technology displacement comes in the form of new system architectures. Shifting service provider networks from analog optics to pure digital optics creates a distributed access architecture (DAA). The fiber network now known as a converged interconnect network (CIN) allows for many advantages for future features and performance.

These advantages include features like low-latency functionality for gaming and Mobile (5G) backhaul, better performance to optimize DOCSIS 3.1 for more capacity/speed, and software features for operational simplicity and support.

High speed access via cable or telco has been a long battle that benefits the consumer in the end in the form of competition for better services and/or pricing. The cable industry has re-invented itself many times over with DOCSIS now with version 3.1 deployed and D4.0 on the horizon. We can easily offer > 1 Gbps DS and potentially 100 Mbps or higher US speeds.

The latest "trick up their sleeves" is a new architecture to exploit this even further. By displacing the modulation/demodulation functionality (the physical layer - PHY) remotely, cable operators can offer better reliability, higher quality and in essence, higher speeds.

So, the question, "why R-PHY or DAA in General", will be addressed to provide answers to questions the attendees may not have even known to ask. We look at Remote PHY technology and implementation concerns for CMTS platforms along with benefits of DAA. A new phase is occurring now for outside plant upgrades that convert analog optics to digital optics. There are lessons learned and best practices to make this transition easier, less cumbersome with less customer-generated support cases.

2. Background

Current remote phy devices (RPD) come in node module or shelf designs and support the DOCSIS 3.1 spectrum requirements. The downstream (DS) upper bandedge is 1.218 GHz with an upstream (US) upper bandedge of 204 MHz. The DS lower bandedge will be dictated by the diplex filter and could be hardware or software controlled. Typical US/DS splits will be 42/54, 85/105 (this may actually be 102 to satisfy legacy out-of-band (OOB) settop box DS signaling around 104 MHz), and 204/258 MHz.

DOCSIS 4.0 full duplex (FDX) allows the spectrum band between 108 and 684 MHz to have US and DS overlap, but only for non-primary DS OFDM with US OFDMA signals. D4.0 frequency domain division (FDD), previously known as Extended Spectrum DOCSIS (ESD), is associated with a DS bandedge extension to 1.8 GHz. This also increases US to 684 MHz with an 834 MHz DS start. The US options come in 96 MHz blocks from the D3.1 204 MHz bandedge, but recently one block edge was removed from the specification; 588 MHz.

The physical layer devices (RF signals) from the CMTS will be remotely located from the headend (HE) to the field, whether in a node or some type of shelf design. Also, the fiber will be converted from analog to digital causing all HE signals needing to be created at the remote location. These devices must now generate multiple simultaneous DS signals such as:

- CW carriers for:
 - Leakage test signals for Arcom, Effigis, Comsonics, and Trilithic/Viavi typically at 138 MHz and 612 MHz. Four total carriers where most vendors use two. Trilithic may use three.
 - AGC pilots and alignment tones. Ability for placement on the visual carrier frequency and levels ~6 dB > SC-QAMs.
- DOCSIS SC-QAMs (1.1, 2.0, 3.0) & D3.1 OFDM (multiple 192 MHz blocks)
- MPEG and DVB video
 - Out-of-band (OOB) signaling for legacy STB; 55-1 = Motorola; 55-2 = SA/Cisco

Upstream signals must now be demodulated in this remote location and the US RF spectrum remotely viewed in some fashion. RPDs and D3.1 CMs on the market today support:

- 8 SC-QAMs (ATDMA/TDMA)
- 2 OFDMA 96 MHz blocks

3. Deployment Scenarios

Figure 1 below depicts a typical legacy analog deployment for hybrid fiber/coax (HFC) networks. I-CCAP stands for integrated converged cable access platform where all the RF signals are integrated in the CMTS. Converged means the CMTS can provide DOCSIS and MPEG-2 video signals providing much more savings in the reduction or elimination of video edge-qams.




Figure 1 - Typical Analog Deployment

Figure 2 below represents the conversion of a legacy analog HFC plant into a distributed access architecture (DAA) and specifically converting the analog optics into pure digital optics.



Figure 2 - Digital Fiber (DAA) Node Deployment

Figure 3 below shows another scenario where a DAA can be implemented. This could be a hub site consolidation or even a small HE merged into an existing bigger HE. Digital optics affords us the distance and performance advantages to do this. This scenario shows a hybrid approach where analog HFC is still present, but an entire hub or HE has been eliminated and replaced with signals generated farther away.




Figure 3 - DAA Shelf Deployment

A fourth idea would be RF port or service group (SG) expansion in the HE. If a current CMTS chassis has 56 DS RF connectors, that in essence is 56 SGs. If a node split or addition is planned, then adding an entire new CMTS is tough to justify, more rack space, power, cost, HVAC, etc. Another option could be to convert one of the RF cards to a digital card and attach via fiber to a few RPD shelves creating more SGs and RF connectors.

4. Remote PHY High Level Benefits

There are many benefits to migrating to digital fiber, but we'll start with some high-level ones first.

R-PHY solves power and real estate problems associated with scale and growth. We can scale out vs scaling up. R-PHY moves spectrum creation from the hub to a node and turns a node into mini-hub. Spectrum represents 100% of services delivered and now this spectrum allocation is targeted closer to the customers. All RF coax, combining and splitting is replaced with fiber, switching and routing.

R-PHY changes HFC from analog fiber to digital fiber and into a 10/100 Gbps Ethernet. The CCAP core can move out of a hub to the headend and turn the hub into an Ethernet switching complex and/or allow hub site consolidation. This enables sharing of commercial and residential plants and provides lower power per SG in the hub. The plant can now become a full-service IP network with simpler fiber design rules and lower plant maintenance costs.

Digital optics also have lower optics costs (10G), more SGs per wavelength and more wavelengths per fiber; 40 vs 16 for analog.

All the core assets such as Video, CMTS, EQAM can be virtualized and allows a cloud-native networking approach and true virtualization.

5. Digital Optics Benefits

There are a plethora of specific digital fiber benefits and the focus of this presentation and paper.

Analog fiber is typically referred to as the "Achilles' Heel" of HFC. Link distance (budget) dictates the optical link performance and usually the entire end-to-end performance. Digital fiber provides much better modulation error ratio (MER), which in turn allows higher D3.1 modulation schemes to be applied with higher speeds.

One of the biggest pitfalls of analog fiber is overdriving the laser also known as clipping. US laser clipping is one fallout of this and well known because of the nature of US noise funneling and ingress in the lower spectrum. DS laser clipping is also a concern when we start adding

more channels, more D3.1 OFDM spectrum, emergency alert system (EAS) kicks in, and ingress as well. When we convert to digital fiber, all this concern goes away. There is no laser clipping! Granted, we still need to be aware of analog to digital (A/D) compression, but it's much less systematic than analog laser clipping.

One of the driving forces for capacity is the US and the limited spectrum (typically 42 MHz in North America) we have to allocate for DOCSIS channels. D3.1 supports a 204 MHz upper US bandedge and many systems are looking at this. With the level of US laser clipping we get with 42, 65, or 85 MHz systems, 204 MHz could be a non-starter. Supporting an US of 204 MHz and higher will most likely require digital fiber, in one form or another.

Another great benefit of digital fiber is much longer distances are supported. DOCSIS originally stated .8 msec of delay, which equates to ~100 miles (160 km) of fiber from CMTS to CM. This is possible with EDFA optical amplification, but at the expense of more delay and a performance hit. D3.1 actually dropped the delay to .4 msec or 50 miles end-to-end since almost 99% of fiber nodes are within that realistic distance. Digital fiber transmissions do not degrade in performance like analog (intensity modulated light). It is in essence on/off; 1/0. It's much easier to regenerate a 1! These signals are sent in time vs frequency. More capacity for analog means more spectrum. More capacity for digital means more time/faster links. It's not unfathomable to envision digital links across the country and networks evolving into data/server farms centrally located and feeding remote phy devices across the country thousands of kilometers away. This helps consolidate real estate (hubs).

Another very big advantage/benefit that is often overlooked is cost per wavelength. Digital fiber allows dense wavelength division multiplexing (DWDM) of 40 wavelengths vs only 16 for analog.

Note: This all assumes analog video is retired or some type of overlay is used, which isn't optimum either. Also, US RF testing and spectrum analysis are now required at the RPD, so any test equipment in the HE that requires RF input will be obsolete unless upgraded.

6. Power and Space Savings

A huge benefit of a converged platform along with digital fiber is power and space savings. Converging DOCSIS with video delivery through the same chassis can save lots of power and rack space from video e-qam removal. Removing analog video and migrating to digital fiber eliminates the analog optic transport saving power and rack space as well. Coax goes away, RF combining, and splitting is removed, but optical splitting/combining needs to happen. Some rack space saved is needed now for switches, routers, and possibly timing servers.

Typically, two 13 RU CMTS chassis can be collapsed into one chassis. Less CMTSs = less:

- Power
- HVAC
- Rack space

Below is a Cisco command to view power requirements for different components.

```
• cmts#show environment power
=====
```

R0	FRU Power	709 W	
2	FRU Power	390 W	<<< Regular RF linecard with RF PIC
9	FRU Power	150 W	<<< RPHY specific linecard
9/1	FRU Power	18 W	<<< DPIC with 8 SFP+ (SR)

The linecard in slot 2 is a regular integrated linecard that provides 8 DS RF connectors for 8 SGs and uses 390 W. The linecard in slot 9 has the US and DS phy modules removed to support remote-phy and lists 150 W. The RF physical interface card (PIC) must be exchanged with a digital PIC which could have up to 8 optical transceivers. That PIC adds 18 W for a total savings of $390 - (150 + 18) = 222$ W. A typical fully loaded chassis with 56 SGs drops from ~5 kW to 4 kW.

Note: Even though the PHY is removed, it still has to be accounted for since the service provider pays for outside plant powering. This Phy is moved to the node causing an increase from ~120 W to 160 W.

If Video is incorporated for a converged solution, there will be much more power and rack savings!

7. The Power of Digital and IP

Another benefit of digital fiber is the conundrum we face with splitting/combining with RF. We cannot combine multiple RF upstream legs without interfering and adding noise/ingress. Node splits equal more SGs, which equals more RF connections equaling more CMTSs. This means more RF cabling, rack space, and powering besides cost. Figure 4 below shows a typical analog deployment with the CMTS feeding US and DS independent connections to analog fiber optic transport.




Figure 4 - Typical Analog HFC Deployment

Figure 5 depicts what happens when we do node splits and must add more CMTS chassis to accommodate the extra US connections from the added optical receivers. One could combine the US Rx outputs and share on an existing CMTS US port, but that's like taking 2 steps forward and 1 step back. Why combine noise and ingress?

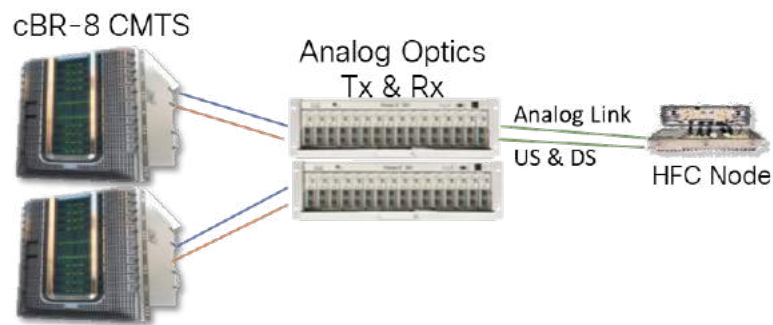


Figure 5 - Node Splitting in RF Domain

Figure 6 below depicts replacing analog optics with digital optics and removing the analog Tx and Rx equipment. Digital optics allows splitting/combining in the time domain with faster digital links. RF does not allow that. If more SGs are needed, they can be deployed in the field without worry of how many RF connections are left on a CMTS in the HE.




Figure 6 - Digital Combining/Splitting

8. Automation and Operational Support

In a large scale RPD deployment, automation is a must. There will be dozens of steps per RPD with anywhere from 50 to 500 RPDs per CMTS core and 100s of cores in typical network or region. There exist some key steps for RPD deployment automation.

1. Initial RPD discovery
2. RPD to MAC resources mapping
3. Config generation and application to CMTS
4. RPD deployment validation
5. Ongoing health monitoring

DAA is the first, critical step in the path to virtualization. Each step works, has value, and is a good investment.

1. Integrated CCAP
2. R-PHY with physical core and physical orchestration
3. vCCAP stand-alone core and R-PHY shelf or node
4. vCCAP in the data center
5. vCCAP with full orchestration
6. vCCAP core with containers and micro-services

9. “Real Life” Testing and Results

DAA with digital fiber links can traverse much longer distances than analog fiber. To prove that a CM could properly connect and provide HSD along with video services, we put together the design shown in Figure 7 below. This was the control test with very short distance.

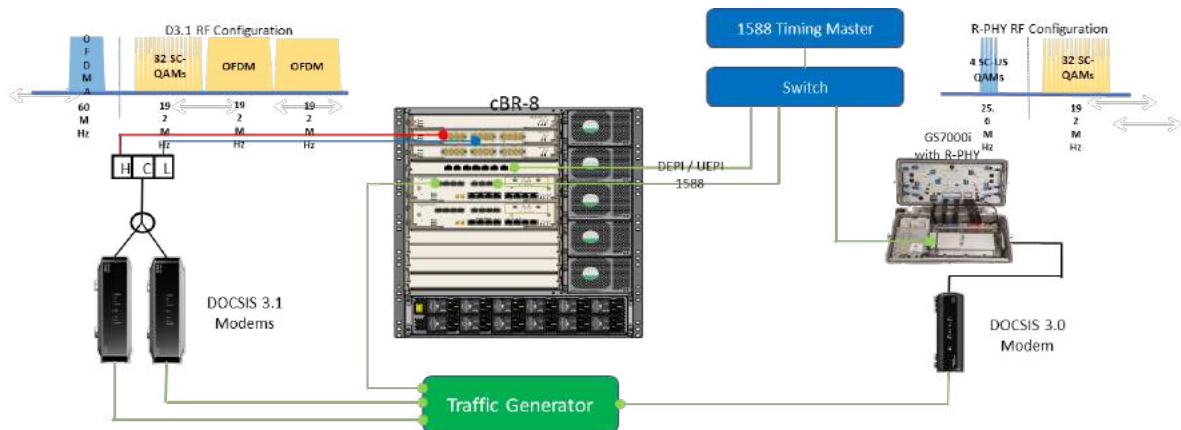


Figure 7 - Lab Testing Diagram

The “real life” distance testing consisted of connecting a CMTS core in RTP, NC with an RPD in Lawrenceville, GA. This is depicted in Figure 8 below. The distance is ~ 350 miles (565 km) at least “as the crow flies”.

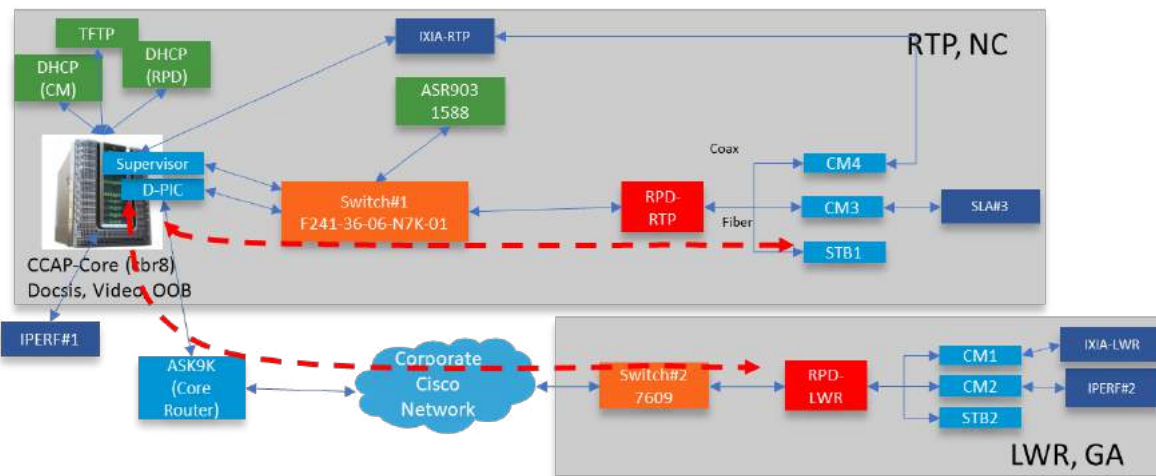


Figure 8 - Real Network Distance Testing Diagram

Running some numbers for velocity of propagation of fiber along with speed of light in a vacuum, we can surmise 2000 km of fiber = ~ 10 ms. This means our fiber distance and delay should only be ~ 3 ms. After some troubleshooting and configuration changes, some observations and major points were made.

- Roundtrip time delay between our CMTS in RTP and the RPD in LWC was 18 ms (avg).

- To our surprise the corporate network required 13 hops to get to our end location. We quickly determined that router and switch delay could be very high, unstable, and unpredictable!
- Using a Linux box for traffic testing required our MTU packet set to 1434 bytes. The MTU across the IT link was set to 1500.
 - When BPI+ was enabled for security, the MTU on Unix boxes changed to 1428 in order to accommodate the extra bytes used in the DOCSIS Extended Header.
- Precision timing protocol (PTP) is used from proper timing between all the components. Stability and reliability are of utmost importance. Properly designing this part of the equation is critical.
- Certain devices in between the CMTS and RPD may require, or be set for, an MTU of 2000 B. Depending on the RPD vendor or firmware, it may be necessary to turn DEPI fragmentation off:
 - `cable depi fragment off`
- Because our distance was > 100 miles and more specifically, > 500 μ s of delay, we had to activate DEPI Latency Measurement (DLM) to properly get the CIN delay added to the Map Advance. If not properly set, the RPD may be fine, but CMs will lose station maintenance and drop offline or not register at all. It is suggested to use DLM with the `measure-only` syntax to track stability of the measurements and if >500 μ s, remove the `measure-only` option.
 - `cmts(config)#cable rpd xxx`
`(config-rpd)#core-interface Tex/1/x`
`(config-rpd-core)#network-delay dlm 1 "measure-only"`
- Because US scheduling is still done at the CMTS for an R-PHY solution, the US request/grant cycle from the DOCSIS protocol can experience delay/latency and affect per-CM US speed. D3.0 with `mtc-mode` and continuous concatenation and fragmentation (CCF) can get around this. The concern is for D1.x and D2.0 CMs if trying to offer > 5 Mbps on the US over such long distances. You could even have a D3.0 CM register in D2.0 mode (US and/or DS) with the same consequences. One option around this is activating a Cisco proprietary feature known as DOCSIS Predictive Scheduler (DPS).
 - `cmts(config-if)#cab upstream ?`
`dps docsis predictive scheduling on one mac-domain`

This feature can even help “speed up” US acks that are required for DS OTT ABR video such as Netflix and Hulu TV.
- Another potential cause for concern with Remote-PHY is the CIN traffic. This part of the network (digital fiber) must not be over utilized or oversubscribed. Traffic engineering is critical and proper quality of service given to critical signaling. Also, what if the CIN is a Metro-Ethernet link not owned by the service provider? All the overhead is being billed with no return on investment. Even if the CIN is owned, overhead is a concern for traffic engineering of the CIN. Some of this overhead traffic is created by DS Map traffic, DS video, NDF/NDR and US triggered spectrum captures.

- SC-QAM primary DSs contain MAPs (500/sec) and that will be related to how many are primary and how many USs are in the mac domain. The mac domain could be a 1x2 architecture causing more USs.
- **Note:** If no data traffic is present during testing, the MAPs are inflated since SC-QAM is MPEG-2 encapsulated leading to 188 B MAPs. Once DOCSIS traffic fills in the 188 B frames, MAPs are closer to 100 B. Also, of note, OFDM primary is not MPEG encapsulated, so 100 B for MAPs is a good number to use. MAPs in DEPI cause UEPI as well, so the CIN MAP overhead will be slightly increased.

Here is an example of potential CIN overhead for a 10Gig link based on a typical deployment.

$$\circ 500 \text{ Maps/s} * 100 \text{ B/Map} * 8 \text{ b/B} * 8 \text{ USs/SG} * (24 \text{ DSs} + 1 \text{ UEPI}) * 2 \text{ RPDs} = 160 \text{ Mbps.}$$

It's possible to get this down to 52 Mbps by implementing D3.1 OFDMA US and removing some US SC-QAM chs along with decreasing the DS primary from 24 to 12.

DOCSIS rate is 38.8 Mbps for 256-QAM Annex B since all layer 1 overhead (FEC and Trellis coding) is done at the phy, so not transferred on the CIN. A raw rate of 42.88 is wrong to use for CIN calculations. That 38.8 may include MAPs as well and will be calculated separately below. The CMTS reports 37.5 Mbps, but actual user-rate is closer to 36 Mbps after subtracting MAPs. One could use 37.5 for the CIN calculation, but 38.8 is probably safer. This is close to 50 Mbps for Annex A Euro-DOCSIS.

Video needs 38.8 Mbps and VoD sends null bytes even if no one is watching. So that traffic has to be accounted for always.

The Map overhead will be worse with 1 ms Maps for LLD, more USs and primary DSs, and multiple RPDs per CIN link. Another overlooked aspect is in European markets that are required, or plan, to keep transmitting FM radio over their cable plant. NDF mode 7 is specified for the digitization of the whole FM band and then sent through the CIN link causing 512 Mbps of overhead! This could be a great time to remove this carriage. We also need to be concerned with US Triggered Spectrum Captures (UTSC) as each could be 50 Mbps. Gathering this information from many RPDs, many ports and continuously could have dire consequences.

10. Future Gazing

“Those who do not learn history are doomed to repeat it”. So, how did we get here? DAA can be many variants such as Remote-PHY, Remote MAC-PHY, Remote Distributed CMTS, Flexible Mac Architecture (FMA) all using a CMTS core (hardware) or virtual /cloud using servers.

There were concerns with Remote MAC-PHY because of cost, complexity, and power requirements. The cable industry settled on a max wattage of 160 W for the node. More complexity in the field (remote) and in many devices brings a level of unease and unknown.

Remote-PHY gained traction and came first because it was simple. There's something to be said for the “KISS” principle. With complex components in a central location and the PHY in field,

it is a simplified design with less chance of issues in the field. Even firmware/software upgrades, bug fixes, and feature additions are easier and more stable.

Knowing that the first iteration of analog to digital migration would probably be a simple node lid changeout with no respacing and 95% of node upgrades would be much less than 50 miles, meant no concern for latency. It's really business as usual (BAU).

With that said, one benefit we listed was hub consolidation, which leads to CIN > 100 miles. Now latency is a concern and it will be necessary to address. This is why we have DLM, DPS and low latency DOCSIS (LLD) as part of the Cablelabs' specification. Part of LLD is the implementation of proactive grant service (PGS) which will provide lower latency for applications that require it like gaming or 5G mobile backhaul. There is even a spec called R-PHY 2.0 that allows for US scheduling in the R-PHY device. So, it is not a full Remote MAC-PHY, but a variant of it.

All this finally brings us to what lies ahead; US scheduling in the RPD explained in the R-PHY 2.0 specification for LLR. Because MAPs are generated in the RPD and the CM US Request and DS Grant are processed in the RPD, there is no CIN contribution to latency. Only the RF portion will contribute and that is typically < 1 mile of coax. This is better for gaming, quicker ping responses and faster US acks for better/smoothier DS TCP, which could be OTT ABR video like Netflix and Hulu TV. This also means that DLM would not be needed.

In addition to less latency, LLR also affords us little to no MAP traffic overhead on the CIN. Having MAPs and the US scheduling in the RPD eliminates reliance on PTP sensitivity or PTP issues in general. These PTP issues could be self-induced, 3rd party devices, link redundancy failovers, etc. Granted, we would still need PTP for Mobile Xhaul and other features, but simple DOCSIS CM stability, station maintenance and T3/T4 timers would be solid.

A few other positives being investigated; possibly implement the US Scheduler in existing deployed RPD hardware with no additional power draw and maybe faster RPD reboots 😊

11. Don't Forget/Neglect the "Little Things"

I would be remiss to not mention the little things that are forgotten or neglected. Always start with Layer 1 (Physical Layer) when troubleshooting and validating connectivity. Just because we are using digital transmissions doesn't mean there aren't requirements for optical power Tx and Rx. We must verify correct fiber (single mode vs multimode) and SFPs. Care must be taken to not crisscross the Tx and Rx. Single mode (yellow jacket) vs multi-mode (orange or blue) is usually easy to identify by color.

SFPs come in different Tx/Rx power. Short reach (SR) are specified for 20 km, while long reach (LR) and ZR are used for 40 & 80 km, respectively. Using ZR SFP optics on a multimode short distance fiber jumper is not going to work properly if at all. Using optical attenuators may be a work-around, but it's still not suggested.

Note: The different values are really based on the Rx sensitivity and the Tx are all the same around 3 dBm. The distance quoted is also based on pure link budget, but in reality, you will lose power from connections, splitters, splices, etc.

Other fiber optic concerns are the actual connections. Does it require LC or SC connectors (node vs optical switch vs RPD, etc.)? Don't forget the patch panels and bulk adapters. Are they angled or flat connections (UPC vs APC)? These are also color coded as black/red for flat and green for angled. Even if a bulk adapter is green, don't assume the connector on the other side of the patch panel is green (angled) as well. Talking about patch panels, how do we justify the cost of many breakout fiber connections with short jumpers and the associated cost of all those SFPs? Here's where copper/10G twinax has been tried or better yet, active optical cable (AOC). Know your options. In regard to breaking out a 100G link to multiple 10G links, QSFPs are available. Newer technology may provide 25G or 40G.

One of the benefits of digital fiber explained earlier was the capability to support DWDM of 40 wavelengths of light vs 16 for analog. This is made possible by optimizing different wavelengths defined in the ITU specification. We must make sure the proper ITU is matched up as a Tx/Rx pair if hard-set, but newer technology may allow adjustable; manual or automatic, at a cost of course.

We may be moving RF out of the HE and getting rid of coax and splitters/combiners, but at a price. We will be replacing RF combining/splitting with routing and switching and fiber connections.

The CIN must be symmetrical for the DS and US (Tx and Rx) which could also be a concern when designing redundancy. An Ethernet ring translates to two different length redundant paths making the RPD DLM and Map Advance adjustment/update critical. The CIN design could also incorporate daisy-chaining to take advantage of a 10G link but only needing maybe 3G at each RPD.

12. Conclusion & Summary

Upgrading or migrating from analog fiber to digital fiber should be an easy sell, but many variables must be taken into consideration in addition to cost. If we look at Remote-PHY in particular, let's separate facts from fiction.

- It's simple and it works
- N+0 **not** required
- Minimum components in RPD yields:
 - Best cost
 - Lowest node and plant power
 - Max SG density for given power budget
 - Best availability
- Same consistent approach for DOCSIS, Video, & OOB
- DAA needed for:
 - US expansion to 204 MHz & beyond
 - Virtualization
 - FDX
- MAC and Scheduler can be scaled as needed since they are central
- Much better visibility of node with RPD vs old transponder technology/nothing
- DOCSIS & video traffic encrypted on fiber

- Centralized software
- Security: CMTS software kept in secure location
- Interoperability
 - Supported by multiple silicon vendors
 - OpenRPD Forum

Remote-PHY can be trialed in targeted areas and doesn't require a complete overhaul of the plant. A great business case could be made for multiple dwelling units (MDU). This could be digital fiber to the MDU with a 1RU RPD shelf feeding the existing HFC plant of the MDU as shown in Figure 9 below. Maybe LLR could be exploited as well for low latency applications.

The shelf can run on 110 Vac or -48 Vdc, so no need for a special quasi-square wave outside plant powering required for a node solution. This could provide 3-6 SGs and provide DRFI DS output levels. The US and DS are independent with support for 204 MHz US and 1.2 GHz DS. Each SG (US and DS pair) could feed a riser with no concern for temperature effects on US attenuation. We could take this a step further and implement Technetix diplexer-less booster amps to "future-proof" the US split. When D4.0 CPE become available, we could upgrade the RPD to higher splits without the need for a plant upgrade assuming taps and passives are specified out to 1.8 GHz.




Figure 9 - DAA for MDU Business Case

Converting analog fiber to digital with all its benefits is a foregone conclusion. Whether that is FTTH with EPON/GPON or DAA with R-PHY, MAC-PHY or some variant, depends on other variables such as cost, time to implement, and comfort level. Either way, the benefits of digital fiber are indisputable and the way of the future.

Abbreviations

ABR	adaptive bit rate
APC	angled physical contact
bps	bits per second
CCAP	converged cable access platform
CM	cable modem
CIN	converged interconnect network
CMTS	cable modem termination system
cnBR	cloud-native broadband router
CPE	customer premise equipment
DAA	distributed access architecture
DEPI	DOCSIS external phy interface

DLM	DEPI latency measurement
DOCSIS	data over cable service interface specification
DPS	DOCSIS predictive scheduling
DRFI	DOCSIS radio frequency interface
DS	downstream
DWDM	dense wavelength division multiplexing
FDX	Full duplex DOCSIS
FEC	forward error correction
FM	frequency modulation
FMA	Flexible MAC-PHY
GHz	gigahertz = 1 billion hertz
HE	headend
HFC	hybrid fiber-coax
Hz	hertz
I-CCAP	integrated converged cable access platform
ISBE	International Society of Broadband Experts
LLD	low latency DOCSIS
LLR	low latency remote phy
MAC	media access control
MDU	multiple dwelling unit
MHz	megahertz = 1 million hertz
MPEG	motion pictures expert group
NDF	narrowband digital forward
NDR	narrowband digital return
OOB	out of band
OTT	over-the-top
PGS	proactive grant service
PHY	physical layer
PTP	precision timing protocol
QAM	quadrature amplitude modulation
RF	radio frequency
RPD	remote phy device
R-PHY	remote phy
RU	rack unit = 1.75 inches
Rx	receive
SCTE	Society of Cable Telecommunications Engineers
SC-QAM	single carrier quadrature amplitude modulation
SFP	shared form factor pluggable
SG	service group
Tx	transmit
UGS	unsolicited grant service
US	upstream
VoD	video on demand

Guide to Managing Cable Network Traffic Congestion

Addressing Capacity and Congestion with Cable Modem Termination System (CMTS) “Tweaks”

A Technical Paper prepared for SCTE•ISBE by

John J. Downey
Sr. CMTS Technical Leader
Cisco Systems
RTP, NC
(919)931-9453
jdowney@cisco.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Top Seven Steps	3
3. CMTS & Cable Interface Suggestions	6
4. VoIP & Service Tiers	11
4.1. Call Signaling Insurance	11
4.2. Service Tiers	12
5. Going Forward and Planning for the Next Inevitable Event	12
Abbreviations.....	Error! Bookmark not defined.

List of Figures

Title	Page Number
Figure 1 – DS Powerboost Example	4

1. Introduction

With unprecedented demand on local broadband networks, operators are under pressure to prepare for congestion and capacity issues before they happen. During the Spring of 2020, most systems experienced the equivalent of one year's growth in 2-3 weeks. The industry "weathered this storm", but we must be prepared for continued demand. Some suggestions will be given to help alleviate and mitigate congestion now and also provide ideas on how to address capacity concerns into the future on cable operator DOCSIS plants.

2. Top Seven Steps

1. Decreasing subscribers per service group (SG) is one of the most obvious choices, but it can be achieved in a multitude of ways.
 - ✓ Increase the amount of SGs and add more licensing/channels.
 - ✓ Decrease service groups (SG) down to one fiber node (FN).
 - ✓ Utilize upstream (US) segmentation. Maybe the node supports multiple downstream (DS) optical receivers and US transmitters or the US utilizes baseband digital reverse (BDR) and can be converted from 1 to 2 US segments.
 - ✓ The last option would be a physical node split, but if utilizing distributed access architectures (DAA) like Remote-PHY, possibly a 1x1 RPD can be replaced with a 2x2 or some other variant.
2. Verifying no **uncorrectable** Forward Error Correction (Uncorr FEC) and "clean" plant should be part of everyone's proactive maintenance plans, but sometimes relegated to the bottom of the list.
 - ✓ Uncorr FEC is basically dropped packets. Regardless of modulation error ratio (MER), signal-to-noise ratio (SNR), carrier-to-noise ratio (CNR), correctable FEC (fixed packets); Uncorr FEC is the most important to your end-customer.
 - ✓ **Note:** Uncorr FEC is not only caused by bad plant issues. It could be from bad timing (time offsets and MAP Advance), poor port-to-port isolation leading to signal "bleed-over", modulation profile settings, and a myriad of other contributors.
3. Increasing capacity without physical node splits or SG changes may be the first goal and can be achieved with a few ideas.
 - ✓ Use the highest US and DS modulation along with the largest channel widths as possible. This may have no cost increase since licensing may be based on channel only and not the rate/speed.
 - ✓ Utilize DOCSIS 3.1 where possible. ***Note:** More speed does not necessarily mean less latency! D3.1 US may exhibit even more latency with ping tests.

- This brings up an all too familiar question, “How many D3.1 CPE are needed to justify exchanging ATDMA chs for OFDMA spectrum?” Is it 10%, 25%, 50%, even higher? If 10% of your users are using 85% of your capacity, they either have D3.1 CMs already or you could identify those “heavy users” with subscriber traffic management (STM) and then give them a D3.1 CM. Assuming it’s not that bad and 10% use 50% of your US, that could be the justification to drop 2 ATDMA chs and use that spectrum for more efficient D3.1 OFDMA. There are some advantages to dropping from 4 to 2 ATDMA chs. D3.0 CMs get back 3 dB more max Tx power and less DS overhead in the form of MAPs. Two ATDMA chs give 54 Mbps aggregate US speed and one could offer a 20 Mbps service and lower. All other US offerings >20 Mbps would use D3.1.
- ✓ Allocate more spectrum for high speed data (HSD) services. This may entail using spectrum once thought to be questionable such as: roll-off, known ingress areas like CB, LTE, Aeronautic band, and the very low end of the US spectrum. DS may require you to “steal” from video spectrum. Analog video reclamation should be an easy sell but converting MPEG-2 video to MPEG-4, over the top (OTT) video is the ultimate endgame and takes more consideration and planning.
- ✓ Utilize/exploit “Powerboost”™. More speed could translate to “actual” vs “perceived”. Powerboost is a term/feature (name trademarked by Comcast) used for faster speeds that can affect perceived speed (DS & US). A command may need to be configured on the CMTS for DS Powerboost activation along with a very large DS Max Burst setting, like 50 MB in the cm file.
- ✓ Figure 1 below depicts an example of a D3.1 CM with 510 Mbps max rate, 600 Mbps peak rate, and 70 MB DS max burst. This can achieve approximately 6 seconds of Powerboost.




Figure 1 – DS Powerboost Example

- **Note:** DOCSIS 3.0 CMs support a TLV for per-CM Peak Rates.
- Utilize to alleviate typical 10% over-provisioning, which is usually done to negate differences between layer 2 & 3 speed reporting.

- Can exploit US Max Traffic Burst for US “Powerboost” as well. US Powerboost may help alleviate perception of lower speed. This can also be exploited to apply a peak-rate and less over-provisioning, which is typically done today at 10% just to “ring the bell”.
 - The “jury is still out” whether this can negatively affect OTT video and other adaptive bit rate (ABR) applications with buffer loading.
- ✓ An US Powerboost can be achieved as well. If for example we wanted to provide a 500x50 Mbps offering, most would typically set for 550x55 Mbps. Instead, we could use Powerboost to allow: 500x50 max rate, 550x60 peak rate, and 50x10 MB DS/US max burst. This should provide ~ 8 secs of US and DS Powerboost.

4. Eliminate Overhead

- ✓ More USs in a MAC domain creates more DS MAP overhead at ~.4 Mbps per US. Increasing US channel size and eliminating too many US channels will help. Another way to get less US channels per MAC domain would be to create more MAC domains. This can be achieved easily when the SG consists of two fiber nodes (FNs).
- ✓ Also, moving to every 4th DS as Primary could save 54 Mbps per 24-ch DS bonding group (BG). The trade-off is less aggregate capacity for D2.0 CMs.
- ✓ Remove “stale” service flows. Some inactive VoIP flows may not be torn down according to their T8 timer and can be done automatically with the CMTS command, `cable service flow activity-timeout 300`. Add this CMTS global command so flows with no activity > 300 seconds (5 minutes) are torn down if the CM/eMTA does not do it automatically.
- ✓ Understand WIFI, VPN, etc. encapsulation overhead along with potential “bottlenecks”.
- ✓ US acks are used for DS TCP flows. D3.0 & 3.1 CMs support and have ack suppression on by default helping to alleviate US acks. For very fast DS TCP flows (i.e. 1 Gbps), this helps alleviate US overhead from the acks required (typically 20 Mbps decreases to ~7 Mbps). The pitfall to this is each ack is now more important and ack suppression is not very active or efficient when the DS flow is slow (IE < 10 Mbps). So, having many OTT video sessions simultaneous across one CM may not benefit from ack suppression since each flow is 3-10 Mbps.

5. Control Abusers and Denial of Service (DoS) Attacks

- ✓ Cloning – DMIC, BPI+, “Hotlist”.
 - The CMTS can control the same mac address on a chassis and identify potential cloned devices. It can also control CPE “appearing” behind multiple CMs, but it has not visibility across multiple chassis. Some

external devices may have this visibility like the DHCP server and better suited for cloning identification.

- If a cloned device is identified, you can disable CM ranging and registration by implementing the Cisco “hotlist” command.

```
(config)#cab privacy hotlist ?
cm                Add cm hotlist
manufacturer      Add manufacturer hotlist

(config)#cab privacy hotlist cm ?
H.H.H             CM mac address H.H.H
```

- **Note:** The CM could still be ranging “all the time”, but it will not even show init(r1) on the CMTS. Some could argue that it’s better to let it register and give it a cm file with network access disabled.

✓ Over-Use/Abuse

- Deep Packet Inspection (DPI) can be used to at least identify “heavy” users.
- Subscriber Traffic Management (STM) takes rate limiting and monitoring down to the least common denominator, which is bytes. There is no bias towards ports, applications, etc. It solely looks at total bytes over a certain time frame and can dynamically drop the CM’s QoS to a lower rate for a given time period.

✓ Arp Attacks, IGMP Joins?

- Arp Filters, Access Lists (ACLs), subscriber-based rate limiting (SBRL).

✓ Expiring Certificates

- Allow/Deny Lists. Cablelabs’ certifications are expiring by end of 2020 and some CM certifications could render it unusable.

6. Optimize CMTS Efficiency

✓ Load Balancing

✓ D3.1 Graceful Profile Management & US/DS Resiliency/Partial Mode

7. Implement Cache Servers

- ✓ **Note:** Netflix, YouTube and other OTT video providers may drop video quality to save bandwidth and/or temporarily halt 4K video offerings.

- ✓ Allowing a new gaming version of Fortnite or Call of Duty to be stored on a cache server closer to your end-users mitigates WAN traffic and overload.

3. CMTS & Cable Interface Suggestions

- **CM Insertion Interval** - CM ranging opportunities

- ✓ (config-if)#cab insertion-interval auto 120 1000 or (60 480)

- ✓ The Cisco scheduler has dedicated time every 60 ms for initial maintenance (IM). The number of CMs online and traffic utilization will automatically make the insertion interval change between those two numbers. Verify with the show controller command:
 - ```
cbr8#sh contr c1/0/2 upstream | in Insertion
Ranging Insertion Interval automatic (120 ms)
Ranging Insertion Interval automatic (120 ms)
Ranging Insertion Interval automatic (120 ms)
Ranging Insertion Interval automatic (120 ms)
Ranging Insertion Interval automatic (120 ms)
```
- ✓ **Note:** The value reported could be an average between a changing insertion interval and report a value that is not an increment of 60 as one would expect.
- ✓ It may be worth experimenting with this Insertion Interval. The cBR-8 defaults are 120 1000, but we have had success in the past with the old uBR10K defaults of 60 480. We have also used “fixed” settings (lowest of 100 ms and highest of 2000 ms) to address maintenance windows. The lower number creates more opportunities for CM registration at the expense of user traffic capacity.
- ✓ A new CM will start ranging typically around 6 to 9 dBmV in 3 dB steps until the CMTS “sees” it, which is about -20 dBmV at the CMTS. Once the CMTS “sees” it, the CM will report inti(r1) as they are doing initial maintenance (IM). This is contention time and CMs will back-off when, and if, they collide. Explained below.
- ✓ **Note:** The CM should quickly go from broadcast IM to unicast station maintenance (SM) for final ranging and report init(r2). This does not have to be in 3 dB steps anymore.
- **US Range & Data Backoff & Init Technique** - Used to minimize collisions in the US
  - ✓ 

```
cable upstream x range-backoff 3 6
```

    - This can be experimented with in case CMs collide at init(r1) and have to back-off. The code allows a CM to back-off randomly between  $2^0$  (1) to  $2^3$  (8) insertion intervals (above command) for first collision. Second time collision, randomly back-off between  $2^3$  (8) and  $2^6$  (64) insertion opportunities.
    - In the case of a CMTS reboot, the insertion interval would be the lowest of every 120 msec and a bunch of CMs would be in inti(r1). Assuming collisions are happening, they would back-off randomly that first time between 1-8 opportunities, this means between  $1 \times 120$  and  $8 \times 120 =$  a back-off anywhere from 120 msec to .96 sec.
    - **Tip:** Look for CMs stuck in init(r1) as they could cause issues by “eating up” limited IM opportunities and cause high Uncorr FEC counters.
  - ✓ 

```
cable upstream x data-backoff 3 5
```

- This is for contention Request collision back-off. The good thing about faster US service flow speeds, is usually after an initial contention Request, the subsequent bandwidth (BW) Requests end up being “piggybacked” within the actual data traffic and no more possibility of collisions. We have seen good results lately in some markets using values of 4 and 6.

**Side Note:** More US utilization coupled with applications not using unsolicited grant service (UGS) such as Vonage, Skype, Zoom, Wi-Fi calling and other BE VoIP, will increase the probability of Request collisions. This could also be exacerbated by DS OTT video and its TCP acks that must be sent on the US.

I suspect customers with audio-only will have more contention requests since video would increase the US throughput requirements and piggybacking should occur more often.

**Warning:** These collisions could lead to laser clipping and dropped packets. This is not the case for distributed access architectures (DAA) like remote-PHY since the fiber link is digital and there would be no laser clipping.

The following Cisco CMTS commands can be used to verify BW Requests whether they are contention or piggybacked. It cannot tell when contention requests actually contend/collide. The first one is intended for a specific CM. Refer to your CMTS vendor for similar commands.

```
cbr8#sh int cx/y/z sid n count ver | inc BW
BWReqs {Cont,Pigg,RPoll,Other} : 8306, 3243, 0, 0
```

This second command will show per US.

```
cbr8#sh contr cx/y/z up n | in Request|Bytes
Bandwidth Requests = 2776290
Piggyback Requests = 1077964
Invalid BW Requests= 195
Bytes Requested = 256264277
Bytes Granted = 1626995783
```

If for example 500 homes were in a SG/FN and 10% are doing some sort of teleconferencing and 40% of them are doing audio-only and half of them actually have collisions. This gives  $500 \cdot .1 \cdot .4 \cdot .5 = 10$  potential request collisions.  $10 \cdot \log(10) =$  a 10-dB potential power spike. To add power perfectly, signals need to be the same frequency, amplitude, and phase. At the US laser input, signals will be the same freq and power, but phase is based on timing/distance. CMs have time offsets to keep tight timing alignment, so phase could be aligned as well.

**TIP:** A trait of laser clipping is “seeing” artifacts like second and third order harmonics above the diplex filter region. One way to prove a signal is an artifact is to turn off the original, “real” signal or watch a spectrogram view, which is time in the Z axis. If artifacts disappear the same time signal below 42 MHz disappears or fluctuates, then it’s a high probability that it’s a harmonic or by-product of inter-mixing of signals (heterodyning). Keep in mind that sometime DS signals leak on the US, so it’s actually ingress and not a harmonic. Also look below 5 MHz

and make sure AM or HAM radio is not getting into your node. It's been seen in the past where a node using a special port for power insertion wasn't as efficient as believed for RF choking. Installing a power inserter on an RF leg solved the issue.

Other power spikes could be CMs coming online and ranging. A CM on a low value tap will normally only need to transmit maybe 35 dBmV and if it ranges it could go as high as 57 dBmV. Utilizing flexible solution taps (FST) with built-in EQs helps alleviate this since CMs all transmit between 40-50 dBmV and will not have a large range to ramp up.

**Warning:** There could also be a concern with CMs in the "hotlist" as they will still range. Whether this exacerbates the issue is unknown since they never show init(r1), but they're ramping up on every UCD and trying all day long!

- ✓ `cable upstream ranging-init-technique 2`
  - This cable interface command helps US ranging for D3.0 mtc-mode (US bonding) by eliminating contention ranging on the other USs in the US BG once the first US has ranged. The default is technique 1, which means contention IM. Tech 2 is unicast, so basically SM ranging. There have been issues with some CMs with tech 2 in 3.18 code, but tech 3 or 4 could be tried as well. It also helps with RFoG systems and also for D3.0 DS load balance (LB).
- **Throttle CM Ranging**
  - ✓ `[no] cable throttle-modem init-rate <1-1000> holdoff-time <5-100> flush-rate <100-1000>`
    - Suggested values; 32 CM/s; 45 sec; 300 CM/s
  - ✓ `show cable throttle-modem`
  - ✓ `cable up rate-limit-bwreq exempted-priority <priority>`
- **Prioritize Pre-registration Traffic**
  - ✓ `(config)#cable qos pre-registration us-priority [0-7]`
    - Default of 0, with a suggestion of 6 or 7. During CM registration, a CM first goes through init(r1), which is contention-based ranging. Once the CMTS "hears" it, it goes to init(r2), which is unicast ranging to fine tune the levels and add Pre-EQ. Once these physical layers are complete, the CM state will report int(rc) and it can now be "docsis pinged". The next state is dhcp (init(d)) and the CM must now use actual data transmissions that compete with other CM transmissions. If the dhcp discover is small, the CM could use a short grant and its associated modulation, otherwise it will use a long grant with its modulation. If other CM transmissions are higher than priority 0 and lots of US utilization, then this init(d) state may never have any opportunities to be fulfilled, so set it much higher than 0!
  - ✓ DS – "cable service flow priority" (EDCS-1524683). Contact Cisco TAC and/or CX for more information.



- ✓ **Note:** Setting all BE flows > priority 0 can lead to issues.

- **US Max Power Issues**

- ✓ `cable upstream n power-adjust continue 6`
  - Helps CMs exceeding Max Tx power to stay online.
  - **Note:** A max transmit CM will be commanded to change level every 15-20 seconds during its SM, optimally only once. Some CMs have been observed to go into the fast polling mode (every second) for 5-10 times before moving on. This is a good reason to make sure < 5% of your CMs with a Rx marking of ! are in this state.
  - By increasing the “continue” command to 6 dB, the CM will be permitted to stay online if the CMTS receive level is between -6 dBmV and 0 dBmV. If the level is above -1, you won't see a "!". If the level is below -6 dBmV, the CM will go offline. For systems that still have high-value taps (29 & 26 dB), this helps keep the CM online, but will produce CMs with different CNRs & MERs.
  - **Warning:** Allowing a large `power-adjust continue` to be configured can lead to CMs having a large range to overcome isolation and potentially appear on US ports where they should not! It could also allow CMs located off low value taps to range very high and create intermittent laser clipping.
  - **Note:** If the level of noise on the US is enough to distort the US level being received by the CMTS, then the CM and CMTS will go into “`power-adjust noise`” averaging mode. A “\*” will be displayed next to the receive level in the `show cable modem` command. When this occurs, CMs are polled using a one second interval. By default, the percentage of “noisy” ranging responses that cause a CM to enter “`cable upstream n power-adjust noise`” mode is 30%. This percentage may be increased to alleviate excessive power level adjustments in the presence of noise.
  - The following command can be used to identify CMs in max Tx power, max time offset and noise averaging mode.

```
cbr8#scm | in *||MAC|State
```

| MAC Address    | IP  | I/F       | MAC          | Prim | RxPwr  | Timing |
|----------------|-----|-----------|--------------|------|--------|--------|
|                | Add |           | State        | Sid  | (dBmV) | Offset |
| 38c8.5cb6.63ca | --  | C2/0/2/U1 | online(pt)   | 15   | !0.00  | 1209   |
| 38c8.5c09.42c0 | --  | C2/0/1/UB | w-online(pt) | 1    | -1.00  | !6104  |
| 6477.7d90.4368 | --  | C2/0/7/UB | w-online(pt) | 12   | *0.50  | 1532   |

- ✓ `cable upstream max-channel-power-offset 6`
  - The above command helps D3.0 & D3.1 CMs select the best US BG. When a D3.0 CM registers, it does so on a single channel, a reference channel, and relays its Tx level back to the CMTS. The CMTS can determine if that level will be adequate for multi-ch bonding. The “`power-adjust continue`” range is **NOT** used for this decision.

- **Note:** Cisco has a feature that will drop from 4-ch to 2-ch (if configured) and finally single-ch mtc-mode. This depends on the Tx level supported plus it adds in the max-channel-power-offset calculations. This command has a default of 3 dB, but a value of 6 is recommended. If that level is not adequate for all options, then the CM resets itself. Example; CM ranges on US0 and reports 55 dBmV, CMTS wants to do 4-ch bonding and determines that 64-QAM for 4-ch US bonding has a max output of 51 dBmV + 3 max-ch-offset = 54, so CM drops to 2-ch BG, if configured.
- ✓ Stick with the double minislot from default like we suggest and never quadruple it. If so, more “time on the wire” will be wasted. Dropping it to the default minislot of 1 when using 6.4 MHz ch width will not save anything and could affect US concatenation and per-CM US speed.
- ✓ cable upstream balance-scheduling
  - The US scheduler tends to allocate more minislots in the first US in the US BG if not using this command. This is not “bad” but can affect D2.0 CM US load balancing. This command is not on by default, but highly suggested to implement. Another option would be to assign US0 to your highest, “best” US frequency and the last US in the BG to the lowest, “worst” US frequency.
  - **Warning:** Do not use this command for RFoG (DPON) environments.
- ✓ cable upstream qos fairness
  - Implement the qos fairness cable interface command to help fairly share between D3.0 and D3.1 CMs so one doesn’t “starve out” the other. It’s not on by default and we have seen D3.1 allocated more speed at the expense of D3.0 CMs. The command doesn’t change the cross-bonding functionality. D3.1 CMs still prefer 3.1 spectrum before utilizing 2.0/3.0 spectrum (chs).

## 4. VoIP & Service Tiers

### 4.1. Call Signaling Insurance

- ✓ Utilize non real-time polling service (nRTPS) for call signaling. This allocates non-contention request opportunities to guarantee call signaling during high US congestion. The beauty of nRTPS is it allows contention requests, if available, along with non-contention requests, while RTPS is non-contention only. It also allows a priority to be configured for the flow associated with the nRTPS request. Keep in mind that this flow will use another SID.
- ✓ This would be a good time to re-evaluate the modulation profile used for the A-UGS burst since much more traffic could be created by eMTAs. Using a more robust modulation vs the A-Long burst may not be in our best interest anymore and cause undesirable wastage of time on the wire.

## 4.2. Service Tiers

- ✓ When adding faster service tiers, be sure to delete the old, slower ones. Many people forget to delete obsolete tiers when they migrate to higher tiers.
- ✓ **Warning:** The slow-to-fast ratio should not be more than 1:1000. If it is, the slower rate could constrain the faster rate!
- ✓ Make sure DS call signaling flows utilize an LLQ flow by making sure they use a non-zero max latency value. Then these slow flows will not affect the ratio limit.
- ✓ Look at all the flows forwarded on a Wideband or Integrated interface and verify the highest rate and the lowest rate do not exceed a 1000:1 ratio.
  - Example: if offering a 1 Gbps speed, then the lowest offering should be 1 Mbps and higher.
  - It's also not a good idea to use a minimum guarantee rate for any flows (US or DS). Dynamic QoS flows like UGS are fine.

## 5. Going Forward and Planning for the Next Inevitable Event

- Implement a subscriber-based subscription model for quick activation of more channels/capacity.
- Have segmentable nodes for future segmentation and quick activation.
- Implement DAA for better performance, complementary to D3.1, and a pathway to Cloud.
- When available, implement D4.0 Low Latency DOCSIS (LLD) features.
  - ✓ One of those features is Proactive Grant Service (PGS).
  - ✓ Cisco has its own feature called DOCSIS Predictive Scheduler (DPS).
    - `(config-if)#cab upstream dps`
    - Helps US latency in long CIN delay DAA, lowers latency for DS TCP.
    - **Note:** Intel/TI Puma 5 CMs don't seem to benefit.

## Abbreviations

|      |                                 |
|------|---------------------------------|
| ABR  | adaptive bit rate               |
| APC  | angled physical contact         |
| bps  | bits per second                 |
| CCAP | converged cable access platform |
| CM   | cable modem                     |
| CIN  | converged interconnect network  |
| CMTS | cable modem termination system  |
| cnBR | cloud-native broadband router   |
| CPE  | customer premise equipment      |
| DAA  | distributed access architecture |
| DEPI | DOCSIS external phy interface   |
| DLM  | DEPI latency measurement        |

|        |                                                 |
|--------|-------------------------------------------------|
| DOCSIS | data over cable service interface specification |
| DPS    | DOCSIS predictive scheduling                    |
| DRFI   | DOCSIS radio frequency interface                |
| DS     | downstream                                      |
| DWDM   | dense wavelength division multiplexing          |
| FDX    | Full duplex DOCSIS                              |
| FEC    | forward error correction                        |
| FM     | frequency modulation                            |
| FMA    | Flexible MAC-PHY                                |
| GHz    | gigahertz = 1 billion hertz                     |
| HE     | headend                                         |
| HFC    | hybrid fiber-coax                               |
| Hz     | hertz                                           |
| I-CCAP | integrated converged cable access platform      |
| ISBE   | International Society of Broadband Experts      |
| LLD    | low latency DOCSIS                              |
| LLR    | low latency remote phy                          |
| MAC    | media access control                            |
| MDU    | multiple dwelling unit                          |
| MHz    | megahertz = 1 million hertz                     |
| MPEG   | motion pictures expert group                    |
| NDF    | narrowband digital forward                      |
| NDR    | narrowband digital return                       |
| OOB    | out of band                                     |
| OTT    | over-the-top                                    |
| PGS    | proactive grant service                         |
| PHY    | physical layer                                  |
| PTP    | precision timing protocol                       |
| QAM    | quadrature amplitude modulation                 |
| RF     | radio frequency                                 |
| RPD    | remote phy device                               |
| R-PHY  | remote phy                                      |
| RU     | rack unit = 1.75 inches                         |
| Rx     | receive                                         |
| SCTE   | Society of Cable Telecommunications Engineers   |
| SC-QAM | single carrier quadrature amplitude modulation  |
| SFP    | shared form factor pluggable                    |
| SG     | service group                                   |
| Tx     | transmit                                        |
| UGS    | unsolicited grant service                       |
| US     | upstream                                        |
| VoD    | video on demand                                 |

# **The Business Case for Aging in Place with Cable Operators**

A Technical Paper prepared for SCTE•ISBE by

**Ian Wheelock**

Engineering Fellow  
CommScope  
ian.wheelock@comscope.com

**Charles Cheevers**

CTO Home Network Solutions  
CommScope  
Charles.cheevers@comscope.com

**Dr. Sudheer Dharanikota**

Managing Director  
Duke Tech Solutions Inc.  
111 Fieldbrook Ct., Cary, NC 27519  
+1-(919)-961-6175  
sudheer@duketechsolutions.com

**Ayarah Dharanikota**

Business Analyst  
Duke Tech Solutions Inc.  
111 Fieldbrook Ct., Cary, NC 27519  
+1-(919)-376-5657  
ayarah.dharanikota@duketechsolutions.com

# Table of Contents

| Title                                                               | Page Number |
|---------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                | 3           |
| 2. Understanding the Aging in Place market opportunity.....         | 3           |
| 3. Connectivity a major element of the Aging in Place solution..... | 17          |
| 4. The Cable Operators solution play for Aging in Place.....        | 25          |
| 5. Business case with estimates of tiered service prices.....       | 26          |
| 6. Conclusions .....                                                | 31          |
| Appendix A: Some solutions operators can offer for AIP Homes .....  | 32          |

## List of Figures

| Title                                                                                            | Page Number |
|--------------------------------------------------------------------------------------------------|-------------|
| Figure 1 US Census 2017 - 2060 population projection.....                                        | 3           |
| Figure 2 2016 healthcare spend from elderly based on their insurance and condition.....          | 4           |
| Figure 3 2019 elderly living by census.gov.....                                                  | 4           |
| Figure 4 Targeted elderly population who prefer to age in place.....                             | 5           |
| Figure 5 Yearly income, in 2018, of Aging in Place individuals .....                             | 6           |
| Figure 6 2019 US monthly median cost of elderly care solutions .....                             | 6           |
| Figure 7 Prevalence of caregiving by age of recipient, 2020 compared to 2015 from AARP .....     | 9           |
| Figure 8 Average number of hours per week spent by caregivers (from AARP).....                   | 9           |
| Figure 9 Caregiver distance from Care Recipient (from AARP) .....                                | 10          |
| Figure 10 Some of the example technological solutions required for the aging population .....    | 12          |
| Figure 11 Simple technology solutions that can assist aging in place .....                       | 13          |
| Figure 12 The two potential solution architectures for AIP services from Cable Operator.....     | 15          |
| Figure 13 Potential 3-way solution set for Aging in Place and Medical intervention.....          | 16          |
| Figure 14 Key Elements of the Connectivity Solutions for Cable Operators in AIP value chain..... | 17          |
| Figure 15 US - 2000-2019 Adults connected to the Internet by age.....                            | 17          |
| Figure 16 The Simple Phases of Aging in Place decline.....                                       | 19          |
| Figure 17 Cable Operator Connectivity can provide the basis for Monitoring Services.....         | 19          |
| Figure 18 The five classes of device types that make up the AIP telemetry solution .....         | 20          |
| Figure 19 The simple monitoring flow diagram for AIP.....                                        | 23          |
| Figure 20 The 5 key tenets of simple AIP solution set for Cable Operators.....                   | 24          |
| Figure 21 Eleven key service offerings from a Cable Operator to support AIP Home.....            | 26          |
| Figure 22 Probability of dying between X and X+1 years.....                                      | 27          |

# 1. Introduction


2020 has been a landmark year for the human race. Never before has the human race been equipped with the technical capability to be able to handle dissecting data as well as innovation in finding cures and prevention to something like the coronavirus strains that now threaten the fabric of how we behave as a society. We are all locked in our homes and trusted areas of interaction with our fellow humans. No sector of the human race has been more affected by the pandemic than the elderly population. The virus itself affects those with compromised or weaker immune systems more often than not. This not only brings our elderly population into the most vulnerable group but also now making elderly care homes possibly the most likely place to catch and spread a virus and contagion.

Even before our pandemic times, it was clear that Aging in Place (AIP) was a key area for new opportunities for Cable Operators to tap into. It's not only a huge improvement change for elderly lives but also a new source of high margin revenue for the Cable Operators. Lockdown at home has given everyone a taste of what it is like spending more time in your home and in particular the role of connectivity in our digital work and social lives. One could almost claim that the foundation pillar of AIP is connectivity. Thus making it obvious that the Service Provider is in a unique position to open up the floodgates on a new defacto model for living out your life in your own home.

Now is the time for Cable Operators to go beyond the triple play and quadruple-play and add high-value connectivity-based services for the AIP cycle of the connected home life. The following sections of this paper will take the reader through the opportunity and some of the key tenets of a Cable Operator led AIP solution. It will also hopefully open up discussion on the key decision points and also the inertia elements for the operator to pivot into this space. This paper will focus specifically on AIP at home rather than in dedicated living communities which can share a lot of the technology and approaches similar to normal residential housing AIP. It will also highlight the simple approach to AIP with the simple tenets of:

- Technology assist for aging in place
- Connectivity
- Communication
- Data Analytics

With a practical focus on the simple assists for Aging in Place, the Cable Operator can do well in the high-value service that will emerge to keep people at home longer.




**Figure 1 US Census 2017 - 2060 population projection**

## 2. Understanding the Aging in Place market opportunity

US aging population percentage is growing. Based on<sup>1</sup> the US Census, as shown in Figure 1, the aging population that is above 65 years old, is nearly going to double from 49M in 2017 to 95M in 2060. This is attributed to the reduction in mortality rate. This growing elderly population is going to significantly increase healthcare and in general, their lifestyle-related spending.

<sup>1</sup> United States Census Bureau, *2017 National Population Projections Tables: Main Series*, available [here](#)




**Figure 2 2016 healthcare spend from elderly based on their insurance and condition**

As shown in Figure 2 from VizHub<sup>2</sup> Out of \$2.7Trillion spent on healthcare in 2016 in the US, ~35% is spent on the above 65 years age group. This amounts to ~\$1Trillion on the needs of an aging population. Based on the yearly healthcare cost allocation from National Health Accounts<sup>3</sup> the overall healthcare spending reached \$3.6Trillion in 2018. If the same trend continues, by the year 2028, the US will be spending ~\$1.6T on the elderly population (as presented in the insert 2018 & 2028 US Healthcare Spend and Highlevel AIP Opportunity). This is the basic spend on the elderly with the status quo of support. Through innovations, different industries are trying to address the needs of the elderly. Some of these healthcare spendings can be

used for their lifestyle changes that potentially can reduce the overall cost.

———— The aging population will be more technical savvy → Open to newer business models



**Figure 3 2019 elderly living by census.gov**

The 65-year-old today entering their AIP journey is more technically savvy and has lived most of their life with the Internet (introduced 30 years back), laptops (30 yrs.), smartphones (13 yrs.), and devices such as Alexa (6 yrs.). So, the 65-year-old today is in reasonable shape to leverage and drive technology themselves for their AIP independence. The 85-year-old today in their AIP journey has probably struggled in their use of technology having missed the key events above as part of their working and earlier life. The tech-savviness along with the serious inclinations to staying independent, as shown in Figure 3, 42.1% amounting to 14.4M<sup>4</sup> households (Not including elders couples that are living alone themselves) in the US fall into this target market. Historically the homes are growing at 1.013 times year over year. This will lead to roughly 16M homes passed by 2028. This is expected to grow to ~30M by 2060. Like any solution a

<sup>2</sup> Tracking personal health care spending in the US | Viz Hub, available [here](#)

<sup>3</sup> National Health Expenditures 2018 Highlights, available [here](#)

<sup>4</sup> United States Census Bureau, *The Older Population in the United States: 2019*, available [here](#)



connectivity and technological-based solution to help people remain in their homes as long as possible targets the majority of people who follow the pyramid sections, as shown in Figure 4, below where the person can remain in their home with technical support even to the point where they need regular help from external sources and even if they need full-time nursing care. For issues of the mind and brain and Chronic health conditions, these typically have to be handled in



**Figure 4 Targeted elderly population who prefer to age in place**


Nursing homes or specialized assisted living communities with onsite medical resources.

For this Market research analysis we can make some coarse assumptions to develop a best-case opportunity size, as follows:


Assuming in 2020 ~25% of those 65+ living on their own currently need the help of technology (such as for their self-care, Independent living, and Ambulatory problems) to improve quality of life and the ability to remain at home, the target customer base would be ~14M people. By 2060 this will reach ~24M. 25% of the 14M can immediately benefit from technology assist solutions.

In 2020 about 8% of the population on a Cable Operators network are over 65 and of this about 25% are living on their own. So a *hypothetical operator with 2M subscribers* has an addressable market of 160K subs above 65 of which 45K are living on their own, which will increase to 75K by 2060. This group will be the most motivated to adopt based on the price points compared to the assisted living options.

### 2018 & 2028 US Healthcare Spend and Highlevel AIP Opportunity



NHE<sup>3</sup> provides a detailed breakdown of healthcare actual spend of \$3.65T in 2018 and a projected spend of \$6.2T by 2028, as presented above (numbers are in billions). Out of this overall spend \$2.6T in 2018 and \$4.5T can be addressed by the initiatives that are part of the Aging In Place initiatives. Note that this opportunity sizing includes all age groups. If we assume 35% of this is spent on the elderly (as shown in Figure 2), it would be approximately \$1T in 2018 growing up to \$1.6T by 2028.



**Figure 5 Yearly income, in 2018, of Aging in Place individuals**

Now let us see if we can figure out some potential indicators as to how much people are willing to pay, should pay, or can pay for AIP based technology solutions. As shown in Figure 5, the non-family households (representative of the Aging in Place individuals) income is skewed towards less than \$50K per year<sup>4</sup>. People age 80 and over who live alone had a median income of \$22K in 2018 compared to \$52K for married couples.

One out of five older adults have income from earnings. In 2018, the median income of the four-fifths of people age 65 and older who are fully retired was \$20,440. The amounts were similar

among all older age groups. The monthly income, with no earnings, is about \$1,700 so you will see that all of the monthly costs of care are more than the earnings of a lot of the US senior citizens. 2017 Social Security Bulletin<sup>5</sup> report that roughly half of the aged population live in households that receive at least 50 percent of total family income from Social Security and about one-quarter of the aged live in households that receive at least 90 percent of family income from Social Security. In 2018, the yearly average assisted living cost<sup>6</sup> is \$45K ranging from \$36K minimum in South Carolina to \$72K in Delaware. Nursing home costs run from average in 2019 annually of \$89K for a semi-private room to \$100K for a private room. Things have increased since 2016 and are still increasing as the aging population competes for Nursing Home beds and are on track for 2028 to be \$120K for a semi-private

### Monthly Median Costs: National (2019)

| In-Home Care                    | Community and Assisted Living         | Nursing Home Facility          |
|---------------------------------|---------------------------------------|--------------------------------|
| Homemaker Services <sup>1</sup> | Adult Day Health Care <sup>2</sup>    | Semi-Private Room <sup>2</sup> |
| Home Health Aide <sup>1</sup>   | Assisted Living Facility <sup>3</sup> | Private Room <sup>2</sup>      |
| \$4,290                         | \$1,625                               | \$7,513                        |
| \$4,385                         | \$4,051                               | \$8,517                        |

**Figure 6 2019 US monthly median cost of elderly care solutions**

nursing home room to \$135K for a private room. As usual, things vary by state from the highest Alaska at \$29K for a private room to the lowest Oklahoma at \$5K per month. The 2019 Monthly median cost for care in each of the categories of In Homecare, Community and Assisted Living, and Nursing home care are presented in Figure 6. These numbers don't show the additional costs for memory care or for severe disabilities which are much more expensive. Traditional home health care aides assist seniors with daily activities of living, light housekeeping, offer medication reminders, and serve as companions. Their wages average \$20.50 per hour, \$164 per day, \$5K per month, and \$59K annually. *It is this cost and service that we are targeting for technology assist care to minimize the number of hours, days of in-home services, or aides to help.* Skilled nursing care typically involves services similar to home health aides, but providers are trained and certified nurses or therapists who are able to offer additional care such as medication administration, wound and injury care, and various types of therapy. Skilled care averages \$220 daily, \$6.6K monthly, and approximately \$79K per year. *We also want to try and use technology to*

<sup>5</sup> The Importance of Social Security Benefits to the Income of the Aged Population, available [here](#)

<sup>6</sup> Nursing Home Costs, available [here](#)

*minimize the visits and hours spent in the home of the home health aides and other medical support.* GoInvo<sup>7</sup> estimated in 2019 that ~\$252B of \$3.5T total healthcare costs are spent in Home Health Care and Nursing Care. Compare this with the combined revenue of all US Cable Companies of \$83B in 2018.

A couple of other discussion points to consider when understanding the Aging in Place dilemma. In particular the role of family careers for both elderly and in particular single elderly parents. Some of the key premises of *aging in place business opportunities with technology solutions are to enable the parent's children and relatives to provide a more effective caring solution* for not only their parent(s) but themselves. The problem of children caring for their parents is a complex one and has many parameters from

- different cultural views of living with parents
- the stress and busyness of the lives with their own kids – not having enough time for their parents
- the hassle factor of looking after their parents vs it being a vocation to do it
- the guilt of kids not doing enough for their parents
- the guilt on parents having to rely on their kids for support
- the economics for care – if the parent(s) can't afford, can the kids pay for the Help services
- Proximity and closeness of the kids or family members to the Aging in Place parent(s)

---

<sup>7</sup> GoInvo, *Where your health dollars go*, available [here](#)

And lots of other factors that many readers will be familiar with the 2020 AARP report<sup>8</sup> provides a very

### **A note on taxation and Medicare coverage for Aging in Place**

In most cases out of pocket nursing home costs are generally tax deductible under itemized medical expenses. If yourself, your parent, spouse or another legitimate dependent is in nursing care primarily for medical care, then expenses related to medical care, lodging and meals are deductible. However, seniors in nursing homes for personal reasons rather than medical, will only be allowed to deduct costs associated with actual medical care, but not meals and boarding costs. Medicare will cover skilled nursing care expenses in very specific situations and is not designed to pay for nursing home or custodial care costs long term. One such situation is when a senior has been hospitalized and released, but still requires a bit of specialized care. Medicare will help pay for short terms stays in nursing homes if they:

- Were admitted to the hospital for a minimum of three days as an inpatient
- Have been admitted to a Medicare certified facility within 30 days of the hospital stay
- Need skilled care like physical therapy, speech therapy, and other types of rehabilitation

Those who meet all of these conditions under original Medicare will qualify for assistance as follows:

- Up to 20 days of nursing care is 100% covered by Medicare
- After day 21 and up to day 100, patients will pay a co-pay that averages \$170.50.
- After 100 days, all Medicare coverage ends and all payments are the patient's responsibility.

**Medicaid Coverage for Nursing Home Care.** This coverage assists individuals with many types of medical care including doctors visits, hospital stays and long-term care services such as those received in a skilled nursing facility. Often, this program covers 100% of these costs, but there may be co-payments for certain beneficiaries. For those who qualify for Medicaid, this is the best choice for nursing care coverage. Your Home does not count against Medicaid if someone else is living there or its less than \$595,000 (or \$893,000 in some states). Medicaid may put a lien on your house to cover expenses on your death.

detailed analysis of the impacts on the extended family due to elderly caregiving. They found that more than 1 in 5 Americans (21.3 percent) are caregivers, having provided care to an adult or child with special needs at some time in the past 12 months. As shown in Figure 7, This totals an estimated 53.0 million adults in the United States, up from the estimated 43.5 million caregivers in 2015. When looking at caregivers for adults only, the prevalence of caregiving has risen from 16.6 percent in 2015 to 19.2 percent in 2020. They also highlight *“Unpaid caregiving is increasing in prevalence and the U.S. population continues to age and live longer with more complex and chronic conditions. Caregivers feel the push and pull of providing care on their time, their financial well-being, their health, their family, their work, and their own personal well-being. They may find themselves in need of information, resources, benefits, or programs—but these things are often difficult to find or access, or too expensive to afford. Unpaid caregivers are serving as a core piece of the health and long-term services and supports (LTSS) systems, as well as the main source for long-term care for adults living at home and in the community.”*

-

---


<sup>8</sup> AARP, *Caregiving in the US*, available [here](#)

|                                            | 2020<br>Prevalence | Estimated<br>Number of<br>U.S. Adults<br>Who Are<br>Caregivers | 2015<br>Prevalence | Estimated<br>Number of<br>U.S. Adults<br>Who Are<br>Caregivers |
|--------------------------------------------|--------------------|----------------------------------------------------------------|--------------------|----------------------------------------------------------------|
| <b>Overall</b>                             | 21.3%*             | 53.0 million                                                   | 18.2%              | 43.5 million                                                   |
| <b>Caregivers of recipients ages 0–17*</b> | 5.7%*              | 14.1 million                                                   | 4.3%               | 10.2 million                                                   |
| <b>Caregivers of recipients ages 18+</b>   | 19.2%              | 47.9 million                                                   | 16.6%              | 39.8 million                                                   |
| <b>Caregivers of recipients ages 18–49</b> | 2.5%*              | 6.1 million                                                    | 2.3%               | 5.6 million                                                    |
| <b>Caregivers of recipients ages 50+</b>   | 16.8%              | 41.8 million                                                   | 14.3%              | 34.2 million                                                   |

\* Significantly higher than in 2015.


**Figure 7 Prevalence of caregiving by age of recipient, 2020 compared to 2015 from AARP**

As shown in Figure 8, family caregivers<sup>8</sup> spend ~24 hours per week on their loved ones while 21% of caregivers are providing 41+ hours of care. The solution potential is to appeal to these family caregivers to give them the tools to have a more balanced life themselves and be able to use technology to be able to decrease the amount of active time that they spend with the parent(s) they are providing care.



**Figure 8 Average number of hours per week spent by caregivers (from AARP)**

AARP analysis<sup>8</sup> also identified that, as shown in Figure 9, ~76% of the caregivers are within 20 mins from the care recipient and ~11% of them are more than an hour away. This also provides a good case for the monitoring solutions that can be offered for the AIP population.



**Figure 9 Caregiver distance from Care Recipient (from AARP)**

As shown in Figure 7, an estimated 40M Americans are providing care for older members of their family. These family careers often have severe impacts<sup>8</sup> on their own careers, finances, and ability to also save for their retirement and healthcare. *The technology solutions suggested here in this paper would offer many of these careers more time and the use of technology to be more efficient with their face to face time with their elder parents and in many cases remain remote for their parents for a lot longer in the Aging in place lifecycle.* Some of the aging population needs are well articulated in the National Science and

Technology Council report<sup>9</sup> as extracted into the insert “*Emerging technologies to support an aging population*” below.

### **Emerging technologies to support an aging population**

This report identifies a range of emerging technologies that have significant potential to assist older adults, and it is offered as a guide for both public and private sector research and development (R&D) *to improve the quality of life, enhance individual choice, reduce caregiver stress, and cut healthcare costs*. The Task Force identified six primary functional capabilities as being critical to individuals who wish to maintain their independence as they age and for which technology may have a positive impact.

**1.Key Activities of Independent Living.** Living independently requires the ability to perform of a range of activities that impact our daily lives. Many of these activities can be assisted through technology, including those that support good nutrition, hygiene, and medication management.

**2.Cognition.** Cognitive changes are common during aging, with increasing prevalence at older ages—varying in severity and impact. These changes can affect the ability to live independently as well as personal safety. Technology holds the promise to help older adults monitor changes in their cognition, provide mental training to reduce the impact of these changes, and create systems that assist individuals and families to maintain financial security.

**3.Communication and Social Connectivity.** Older adults may face communication challenges as the result of hearing loss, social isolation, and loneliness, especially in economically distressed and rural communities. Technology can improve hearing and strengthen connections to larger communities.

**4.Personal Mobility.** Mobility is a key factor in successful aging. To live independently, an individual must have the ability to comfortably and safely move around the home and throughout the larger community. Technology can assist older adults in staying mobile and able to safely perform key activities necessary for day-to-day life as well as interact with their communities.

**5.Transportation.** True independence requires mobility outside of the home and neighborhood. Transportation needs and limitations are dictated to an extent by the changes to individual physical and cognitive abilities that come with age. While some older adults remain completely independent and continue to drive without assistance, others may be able to drive but require vehicle modification and/or advanced technologies to assist them while operating a vehicle. New technologies could also help older adults more safely and easily use public transportation.

**6.Access to Healthcare.** Access to healthcare plays a critical role in helping older adults stay active and independent as they age. Activities and strategies that support the maintenance of function and independence with age are multifaceted. Alignment and coordination of these efforts through technology can increase the effectiveness and efficiency of these services.

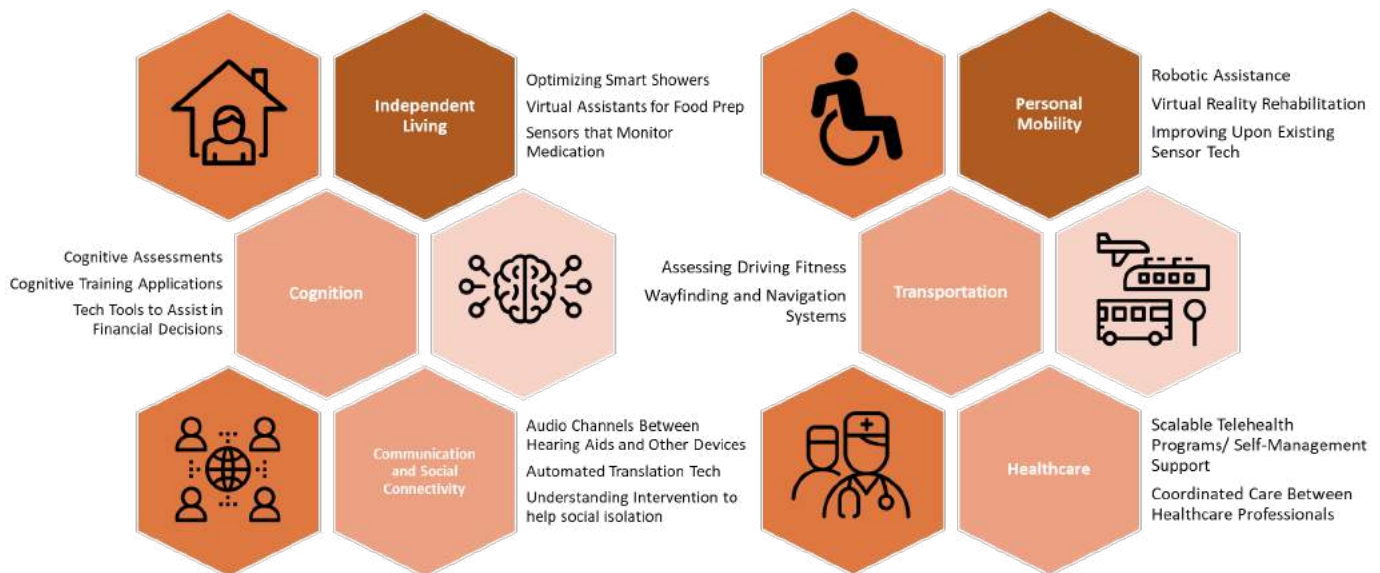
In the process of identifying primary capabilities and focus areas in which technological advances can have a positive impact in enabling older adults to age in place, several areas emerged that are associated with a number of technological solutions and were therefore not specific to individual R&D recommendations. These areas are included in the final section of the report, **Cross-Cutting Themes**.

---

<sup>9</sup> Committee on Technology of the National Science and Technology Council, *Emerging technologies to support an aging population*, March 2019, available [here](#).



As shown above, significant work is in progress to address the aging population needs and one of the top priorities is to provide this to reduce the overall healthcare costs. It is also intuitively understood that aging in place will provide significant cost savings. Note that these costs are not all related to healthcare alone, as explained in the insert “*Emerging technologies to support an aging population*”. Some of the technological areas aging population require assistance is highlighted in Figure 10. Although this gives the depth and breadth of the needs for the aging population, the stepping stone for this is an **AIP Home**. The rest of the paper we focus on this application noting that the opportunity for the Cable operators is a lot more than what is covered in this paper.



**Figure 10 Some of the example technological solutions required for the aging population**

As a first step to the business case, let us understand the current costs or payments by the elderly population. Then analyze this against the target population who can adapt to the elderly-friendly smart homes. In later sections, we analyze some of the solutions, as low hanging fruits, that Cable operators can offer to reduce the health care costs and its potential revenue opportunities.

Figure 6 provides the cost of the elderly population in different stages of support. The following discussion will outline different opportunities for the Cable operators -

- Provide a preventative/pre-emptive technology solution to the 55M maximum elderly homes
- Out of which 3.5M homes that can potentially use the solutions now with specific packages
  - o This opportunity is going to grow to ~6M by 2050 based on population growth estimates
- Offer the technical solutions to reduce the human support for in-home care support costs
- Offer solutions to prolong the need for Assisted Living through different technical solutions
- Offer solutions to enable in-home nursing through technologies rather than Nursing Homes

Before we leave the business case roll up it is also important to consider the cost of insurance typically paid for the Aging in Place journey of your life. Fidelity Insurance<sup>10</sup> estimates that the average US Citizen is going to need \$285K to cover the costs of Health in retirement. Annuity.org<sup>11</sup> shows that an 85-year-

<sup>10</sup> Fidelity Insurance, *How to plan for rising health care costs*, August 2020, available [here](#)

<sup>11</sup> Annuity.org, *Health care costs in retirement*, 2019 survey, available [here](#)



old couple in 2039 is likely to be spending \$34K on their Health Care costs not inclusive of Long Term Care specific costs.

There are many healthcare costs related to technical challenges that can be addressed in an *AIP home*. Note that healthcare is only a small portion of the AIP needs (as highlighted in the *Emerging technologies support for AIP insert*). For example, technology can help with,

- Maintaining contact with family and neighbors to avoid loneliness and depression
- Minimizing the likelihood of accidents and issues in the home
- Getting to someone quickly when falls or other medical issues happen
- (Possibly) Following doctors' orders on diet and medication

A simple representation of some of the technology services that could be implemented to curtail physiological, medical issues as well as improve help time and quality of life for those with deterioration in being able to look after themselves. Figure 11 shows some of the technological services that can be offered to the aging in place population.

|                                     |                                          |
|-------------------------------------|------------------------------------------|
| Panic and Help immediately services | Security to remain alone                 |
| Food and Toilet Monitoring          | Health Eating/Drinking and Digestion     |
| Normal Trend deviation Services     | Changes and Threshold crossing alerts    |
| Location Services                   | Where the AIP person is and trend        |
| Medication control and monitoring   | Deterministic medication remotely        |
| AIP Home as Virtual Room of Carer   | Instant remote access to AIP Environment |
| Wellbeing - Security                | Independent control of home entrants     |
| Wellbeing – Sensor Network          | Health Deterioration alarm               |
| Simple Carer to Caree Communication | Phycological – Injury prevention         |

HIPAA considerations: Medical information is protected by HIPAA<sup>12</sup> (Health Insurance Portability and Accountability Act) Patient Health Information (PHI) handling rules. It is important to stress that the Cable Operator offers service to an *AIP Home* and the person is **NOT** to be a medical service provider. ~~or potentially not even to perform the functions of a Care Portal (Where any telemetry from the AIP Home is stored).~~ It is worth exploring this more as one of the main

**Figure 11 Simple technology solutions that can assist aging in place**

questions in the business case for Aging in Place - the liability and responsibility of Cable operators in *AIP Home* service. The collection of telemetry from the Aging in Place home that collates to provide services to the AIP person is the activity we are targeting for the Cable Operator opportunity. The information in the telemetry can be wide-ranging from –

- Non-PHI related AIP application activities - watching TV, motion derived from Wi-Fi detection
- PHI related information like the collection of BLE based health sensors - thermometers, pulse oximeters, blood pressure cuff
- Not so certain PHI information which on its own is not medical information but can be used to derive health issues and is not subject to HIPAA PHI on its own
  - o the collection of other sensors BLE, ZigBee, Wi-Fi, Z-wave, others that perform other detection sensory functions (room presence, door open/close, fridge open/close, pillbox opened/closed)
  - o the use of cameras for visible monitoring and internal home transparency to approved carers

<sup>12</sup> HHS.gov, *Guidance Regarding Methods for De-identification of Protected Health Information in Accordance with the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule*, Nov 2012, available [here](#)

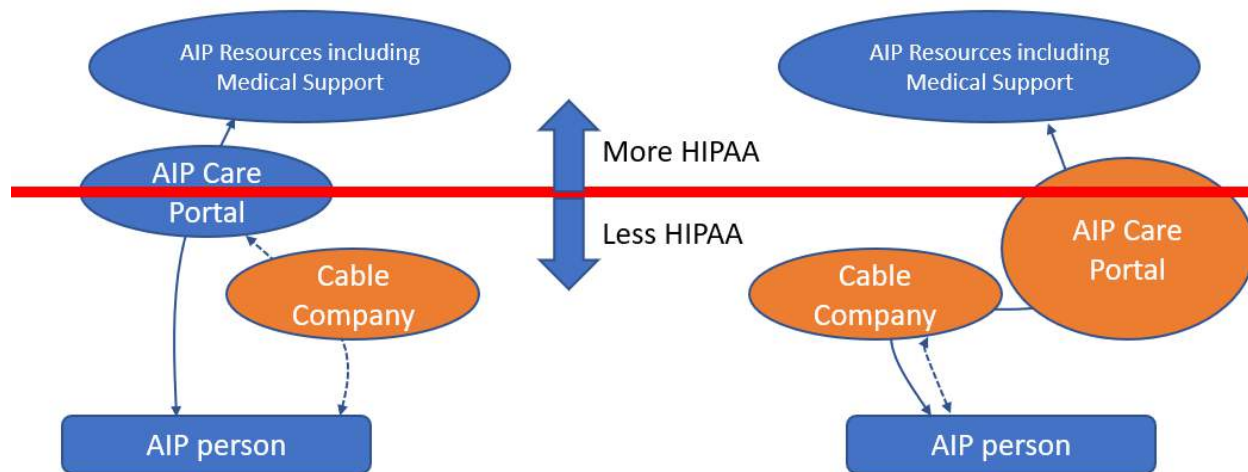
- the use of audio devices to detect cries for help
- the use of panic buttons for emergency help solutions
- Highly PHI associated information such as the specialist internet-connected home medical equipment ranging from pill dispensers to chronic illness support like cloud-controlled and monitored dialysis machines.

These all have varying degrees of issues with privacy, reliability, a consequence of correctness. It is envisaged that most Cable Operators will not provide the AIP care portal solution (for the purposes of this paper's scope – think of it as the interim database of the parametric information collected from the AIP home to perform the AIP service. This database and API interfaces are subject to the rules of privacy for both home and medical (if recorded) information). As shown in Figure 12, the cable operator can use two simple mapping models.

(Model 1) Cable Operator provides **only** a connectivity solution and a set of cloud services to the AIP solution and AIP care portal – forwarding AIP telemetry in a pass-through model to Healthcare Delivery Organizations (HDO's). The client solutions in the AIP solution set are typically supplied by the AIP portal company and in many ways are anonymized to the Cable Operator who may only provide a certain hub and display and audiovisual help services – driven by API's that can be accessed by the AIP solution provider. This access to this cloud, App, and API services can be a 'charge for' service – particularly gaining access to Operator provided device resources in the home. HIPAA compliance is a wide-ranging area which we will outline briefly later in the paper but this solution offering works to abstract the Cable Operator as much as possible from any HIPAA compliance requirements. With the increase in privacy requirements for every Cable Operator and consumer engagement, Cable Operators are already becoming familiar with the security and privacy of data and have implemented solutions themselves and with vendors like CommScope and others to ensure compliance to privacy and security. In this architecture, the Cable Operator provides a business service to

- AIP Care Portal companies who make the compliant linkages to AIP homes and the AIP resources that need to help them like Care Companies and Medical support
- Collects a fee from - AIP person or carers, selected Care Company service provider – allowing them access IoT Hubs, Displays, another home telemetry including voice and audio services.

(Model 2) An argument could be made that the Cable Operator could actually perform the Care Portal function explicitly as an extension of its own Home telemetry databases. The Cable Operator already performs smart home and IoT functions and many of these functions overlap with the same schemes for AIP solutions where some of the IoT devices are recording health information. The Cable Operator can provide the necessary security, privacy, and anonymization features to the northbound AIP resources including medical analysis services. In this scenario, the Cable Operator could expand its Smart Home or Security services to cover some of the functions of response to AIP issues and response requirements. The Cable Company could also partner with AIP solution providers to give them access to their Care Portal solution as part of their service. The Cable Operator would take a percentage of the direct to AIP person service charges for the use of its Aging in Place Care Portal and in-home devices and telemetry services.



**Figure 12 The two potential solution architectures for AIP services from Cable Operator**

Expanding on the issues around liability and having now introduced the HIPAA and PHI words let us spend a few minutes and discuss some of the key areas that typically come up when discussing Health over Cable solutions

- Medical Liability when offering services that can be tagged as pulling health or medical information from the home or devices
- Offering emergency alert solutions like panic buttons and not fulfilling SLA with consumers on the reliability of emergency alert solutions.

As stated earlier – the goal of the Cable Operator Aging in Place solution is to pick the solution set to **minimize** any interpretation or prognosis of medical data and to align an AIP solution more with the Home Security offerings of the Cable Operator – where the key issues are reliability of connection to the cloud – and the ability to hand off the telemetry information to the Security service provider and call-out services.

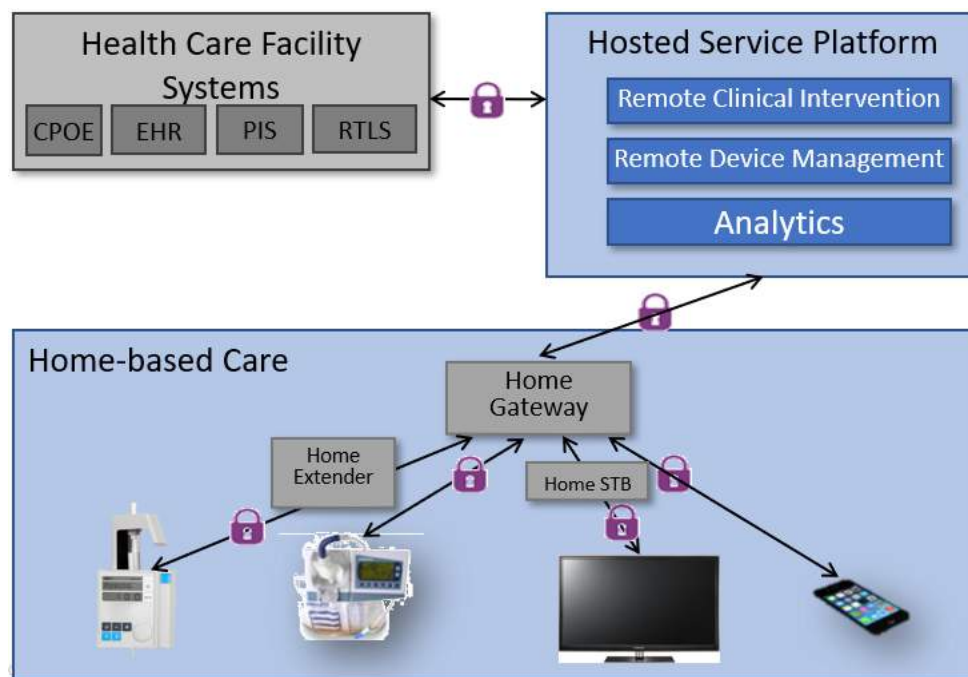
HIPAA (Health Insurance Portability and Accountability Act) is a broad-ranging Act but the key element of it for AIP solutions is to adhere to the HIPAA privacy regulations and provide the defined protection and confidential handling of health information. The National Institute of Standards and Technology (NIST) and the National Cybersecurity Center of Excellence (NCCoE) have been defining standards to make sure that medical devices cannot be compromised, and privacy standards are maintained. There is a new term used on the Internet of Medical Things (IoMT). There are probably two basic levels of worry of exposure

- Informational leaks about the individual's medical conditions, health status, changes in health etc.
- Unauthorized, incorrect, or illegal changes in any medical parameters or status that can inadvertently affect the person or are illegal. This particularly applies to any AIP solutions that have changes made in the conditions in the home like
  - o Changes in recommendations for medication dosage via reminders or other technology assists
  - o Changes in chemical mixes in sophisticated medical devices like infusion pump based delivery systems.

To reemphasize, the cable operator approach to Aging in Place solutions will be to provide the best possible solution to minimize exposure to the deeper medical elements which will typically be offered by partnered solution and the key thing for the Cable Operator solution will be to leverage

- Devices and Hubs to connect to overlaid AIP devices
- Provide reliability of connectivity service
- Provide an AIP Consumer Experience in its own Consumer Experience across its ownership of devices like a TV screen
- Leverage its non-PHI telemetry to create economic telemetry streams when correlated with AIP devices strengthens the Cable Operator value in the AIP value chain

A typical reference architecture for Aging in Place and intervention Telemedicine is shown Figure 13.

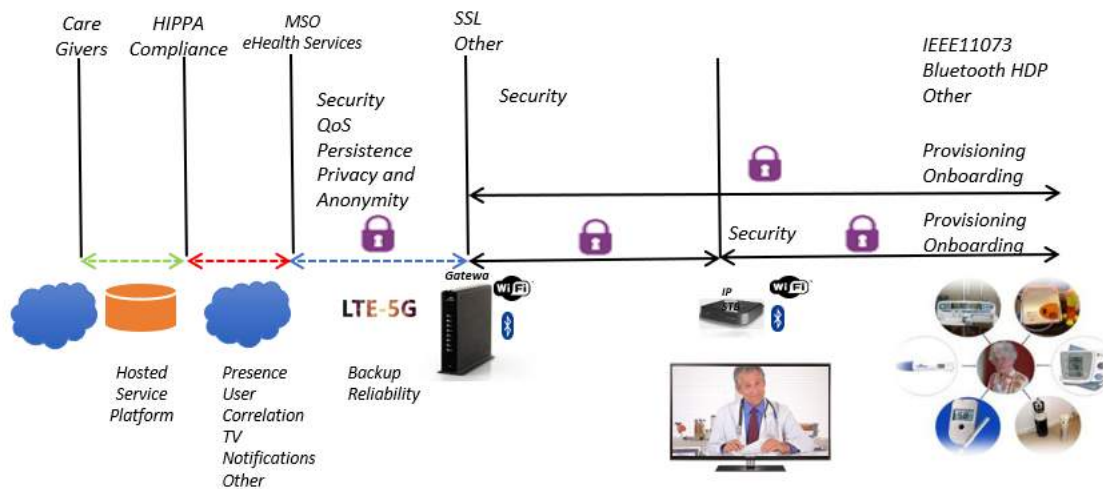


**Figure 13 Potential 3-way solution set for Aging in Place and Medical intervention**

A more detailed but simple illustration of the AIP value process for the Cable Operator architecture in AIP is illustrated in Figure 14. Whatever path the Cable Operator goes in the value chain of Aging in Place they are positioned well to be able to offer the service or partnered services through their

- Ability to market services into the 65+ bracket homes they provide broadband and video services
- Ability to market to the kids and caregivers of an elderly parent(s) who are also in their network
- Potentially also collaborate with each other on AIP services across different partnered Operator Networks. This is the case where the AIP home is in for example Mediacom's network and the Kids and Carers are remote on a Comcast network.

It is this ability to market, fulfil and support the service as well as the integration of the solution into the Connectivity and Video offering (detailed later) that drives a strong case for the Cable Operator to insert in the value chain of Aging in Place.




**Figure 14 Key Elements of the Connectivity Solutions for Cable Operators in AIP value chain**

### 3. Connectivity a major element of the Aging in Place solution

Now, let's start to develop the pillars for the Technology solution and Connectivity solution offerings that the Cable Operator can offer to the Aging in Place home on its network.

The key is using the broadband connectivity provided by the Cable Operator to form the basis of the AIP solution. As shown in Figure 15, the 65+ group is one of the demographics with the lowest internet connections and in 2020 it is ~80% of the population. As part of technology based AIP solutions, there will need to be some new additional connections and the new costs introduced for broadband services. Cable Companies should consider some new reduced rate services for these initial connect applications to onboard the AIP home and then grow the value of the connectivity with the payers of the home connection.



**Figure 15 US - 2000-2019 Adults connected to the Internet by age**

One additional point of note is that there are many people who are for connection to elderly people using technology tools themselves for communication and in some cases do-it-yourself monitoring and even more

advanced health-based monitoring solutions. These solutions typically comprise of – video conferencing capable phones (Apple Facetime, Google Duo) for visual contact, smart speakers (Alexa Dot and Show, Facebook portal, Google mini, Google Next Hub), solutions like Alexa Drop-In – which allows an incoming call to be answered without the user having to pick up, emergency alert pendants, cameras,

some sensors for movement etc. But typically they are not going beyond this small subset of solutions into more complex and linked monitoring and activity promoting solutions.

However, it should serve to highlight the basic importance of connectivity. To create that connection from remote carer to AIP person and to make the visibility into their current status something that can be checked by simply looking at an application on your phone. This is key also to the business proposition of Aging in Place. To be able to

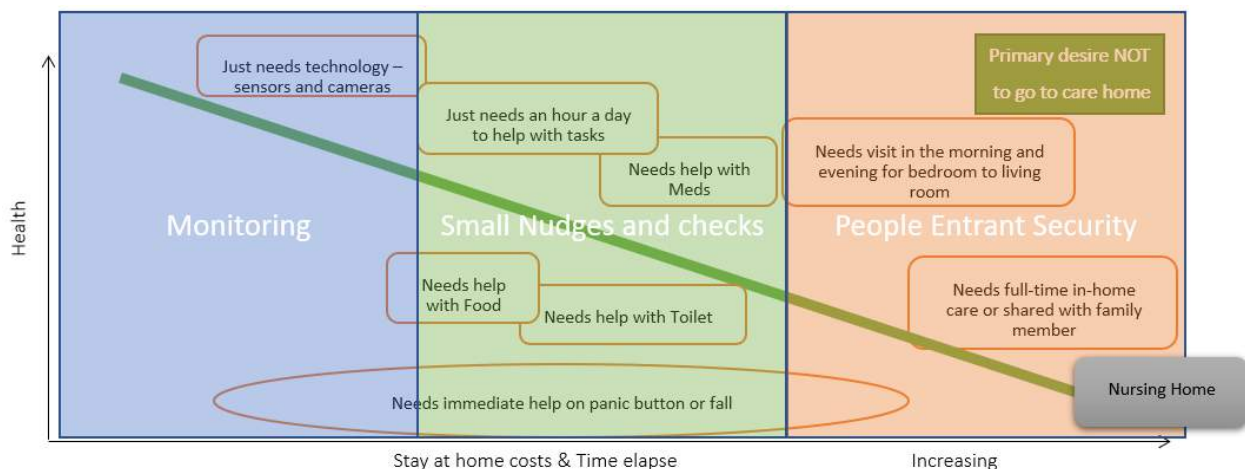
- Make the carer infrastructure remain remote from the house as long as possible
- Only have the carers in the AIP persons home when they desire or is necessary or scheduled
- Minimize the carer travel cost and optimize the efficiency of the time they spend face to face
- Allow multiple carers to share the burden of care with technology to help them all see the daily life metrics of the Aging in place person and spread the response amongst more than one individual for different activities (Remote carers can take a watch on cameras placed in AIP persons home or monitor alerts to alleviate this burden from carers)
- Allow the AIP person to remain in control of their own care needs through initiating help
- Become independent and remove burden that an AIP person may cause on their family
- Spend wisely by the AIP person or their family carers in covering any AIP costs. This sometimes translates to not spending money on even Internet connectivity for the elderly and certainly slower to invest in technology solutions and Do-It-yourself plans.

The last point ‘spend wisely’ is an interesting part of the AIP journey and especially when a decline in ability is slow but steady. Our elderly from 65 onwards, are a generation where saving for ‘a rainy day’ has always been in their minds as well as trying to also leave their kids something to improve their lives. Additionally, when small issues around ease of mobility or getting more chairs bound to creep in, there is not an immediate tendency to start spending money to get help a few times a day to move to different parts of the house. Indeed many or all AIP persons tend to soldier on at tasks that were once simple like dressing in the morning but have gotten much harder due to arthritis or other ailments – and don’t think to look for a morning help service until they reach absolute inability to do the task. People and the Elderly with limited and finite resources will not spend money for help until its usually past the time they needed it. *It is this fear of opening the dwindling bank account or running up debts that many times inhibits an AIP person from using \$20 per hour to help resources to enable them to keep quality of life and continue to function well in their homes. This can be one of the biggest drivers of creating a Cable Operator connectivity led Aging in Place solution that provides elements that always show the cost efficiencies and return on living at home longer.*

The Figure 16 below illustrates a typical three phase potential decline in the Aging in place process. For most people, the journey starts where you remain independent even if afflicted with some health and mobility issues. Most of the solutions for technical support of AIP rely on sound mind and diseases of the mind while they can be still supported with technology often and most of the time involves rely on heavy personal contacts of carers.

As you can see in this very simplified approach to Aging in Place – the first phase of reasonable health is the ‘monitoring phase’. In this phase, connectivity is used to provide remote access for carers to the health of the AIP person and their daily status. *One of the key elements of a Cable Operator Aging in Place offers is to promote the on-ramp of this stage process to the 65-year old that probably feels they don’t need this phase of monitoring until it is required which is sometimes too late for the avoidance of issues.* We will discuss some on-ramp options for Cable Operators later in this section– particularly with the monitoring service tied to existing devices the AIP person has in the home for access to internet and video services.

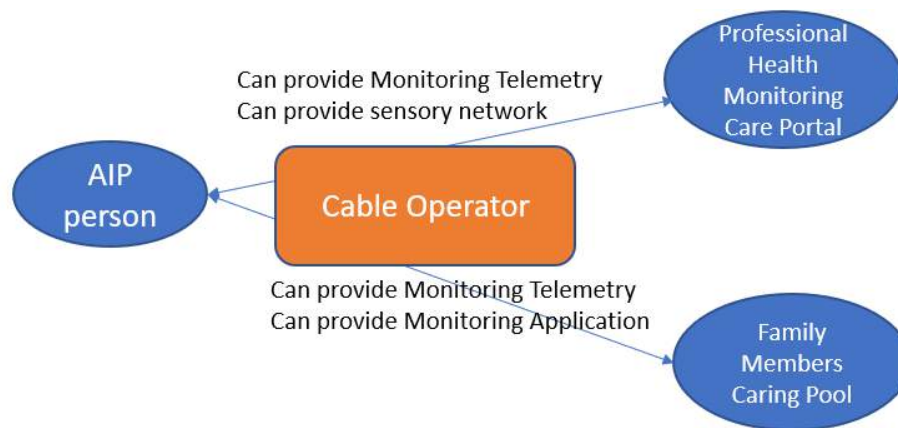




**Figure 16 The Simple Phases of Aging in Place decline**

Monitoring needs to be clarified before we progress. Monitoring, as shown in Figure 17, has usually two aspects when discussed with Aging in Place solutions

- Family carers monitor a parent or parents who are starting to decline a little or through a period where they have to take medication for improvements – *the focus for Cable Operator solution is probably to define a package and market to this audience.*
- Professional monitoring where any home readings go directly to a Care Portal for professional review with an intermediary (doctor or care company) – *this can be done with business development functions directly to the AIP companies in the area.*



**Figure 17 Cable Operator Connectivity can provide the basis for Monitoring Services**

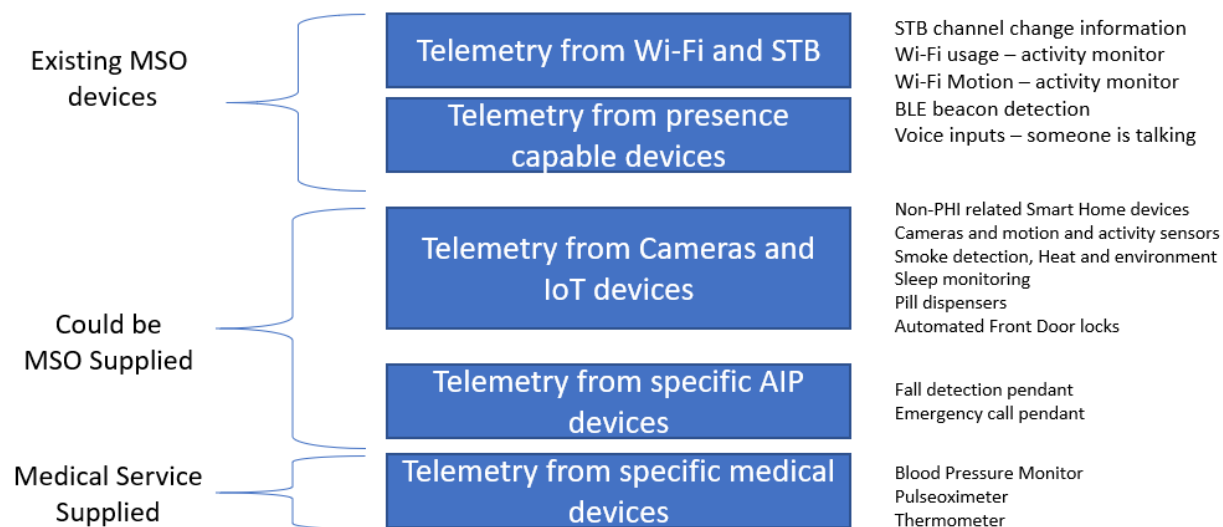
There is an opportunity for the Cable Operator to really own the first and to partner with companies who provide the second option. During this monitoring phase – the focus for a connectivity solution is to be simple. There is a simple hierarchy of monitoring all of which can be offered at different prices and levels of sophistication. The first is the availability of the connection catering to different requirements of reliability and price points.

- High Available connection to the home – typically with a Cable/DOCSIS/Wi-Fi first and backup with LTE or 5G NR (which can be pay as you use basis to make it more cost-effective)

- Leveraging the general DOCSIS broadband connection for the majority of all monitoring solutions and typically of high availability. *Solutions can be added to the Cable Operators' own monitoring of their network to immediately alert any AIP monitoring solutions when the DOCSIS network has dropped – allowing the carers to use other methods of contact (direct calls to LTE/5G based phones or get someone to the AIP person).*
- Leveraging the existing device ecosystem in the home to create a convergent solution offering across the Broadband and Video services which offer reuse opportunities for Aging in Place.
- Leverage the increase in deployment of Smart Speaker assistants that support both Audio and Visual feedback and output.
- Leveraging the increase in deployed IoT radios, in particular, BLE which is becoming the key remote control interface to Set Top Box and Smart Media Devices.
- In particular the key use of the connected Smart Media Device as the connected Hub of Aging in Place services.

Bringing back the connectivity architecture, as shown in Figure 14, the key elements of the connectivity chain, ensuring security on the link, ensuring persistence of data sending, ensuring privacy and anonymity, and ensuring simple onboarding and provisioning of AIP devices/services. All based around the key devices of broadband GW and Wi-Fi, and Smart Media Device (with inclusive BLE and Far-Field Voice and Speaker technology).

Much of the AIP service can be served with this simple and for the most part – a normal home device and connectivity ecosystem. *This is key to the Cable Operator business case in that the investment in AIP capital expenditure is aligned with the general capital expenditure for the new Cable Home inclusive of IoT, Smart Assistants, and the migration of the STB to the Smart Media device.* The SMD will be explored in more detail in the next section of the paper as it is the key device and hub to offer the AIP Cable Operator service. The key areas of transparency in the home and ability to capitalize on both the standard Broadband and Video devices are issued and the general IoT devices supplied as part of smart home and security applications. There is a natural sensory network now emerging in the home – with (i) Wi-Fi (ii) BLE (iii) other IoT radios like Zigbee and Zwave (iv) Voice input – that can provide the foundational pillars of Cable Operators Aging in Place hub and interface to partner AIP services.



**Figure 18 The five classes of device types that make up the AIP telemetry solution**



Probably the most important connectivity and monitoring feature is just the ability to easily contact the AIP home remotely but with the following problems that a better solution than just calling them on the phone has to offer

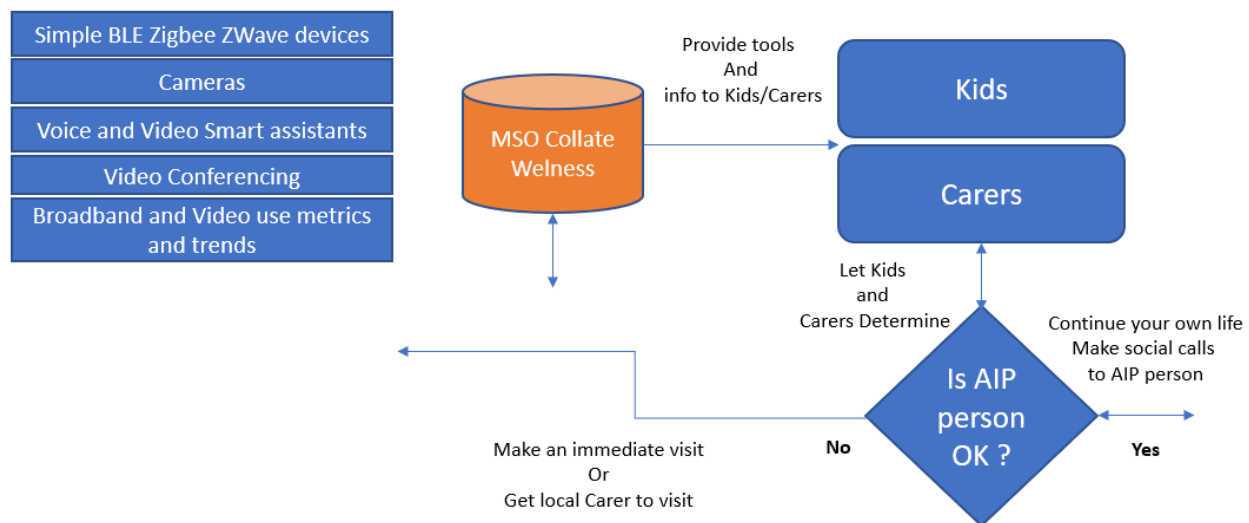
- i. **Video Chat** is always better than just voice for better and more personal interaction as well as the psychological benefits of more intimacy. Additional regular video views of the AIP person helps to do visual checks from carers. *Video transparency of the AIP home is probably the #1 fundamental connectivity feature and it alone could be improved to form a basic package to the AIP home*
  - a. As Elderly people age – the ability to use the Tablet and Smart Phone can diminish. Phones can get lost and require mobility to find them. The use of static Video Conferencing solutions is a key part of the communication reliability to the AIP home. Cable Operator provided Video Conferencing solutions are an important part of the AIP solution. One of the key functions of the evolution of STB to Smart Media device is to support Video Conferencing on a TV – the primary screen that most elderly’s spend a large portion of their day and is present in the rooms they occupy most. The TV itself also has the advantage of size and accessibility features for hard of sight and hearing. Even the humble TV remote can also be a more friendly interface device than a complex smartphone.
- ii. **Use of Voice as an input to the Aging in Place solution.** The advent of Smart Assistant technology has now found a home as part of the Aging in Place solution. Providing a cheap way to have whole-home coverage of a smart speaker and communication solution but most importantly as a way to be able to offer remote intercom support easily. Being able to reliably contact Mam or Dad when they have mobility issues and always be able to connect with them is key to the immediacy of monitoring – feeling that your parent(s) are in the next room and simple check upon them with an intercom solution. Today Alexa Smart Assistants for example offer a drop-in capability where a carer can immediately drop into an AIP home with Alexa’s drop-in enabled and the incoming voice or video call is automatically answered by the AIP Alexa device. This is an enormous advantage in the process of communication with elderly AIP persons with more and more of their time spent in one or 2 sitting locations during each day. Removes the issues with regular phones and smartphones
  - a. Always powered in Smart Speaker – so no battery rundown of mobile phone
  - b. Permanently located in the room and always available – DECT and Cordless phones getting lost
  - c. Does not require the AIP person to get up to find or answer the phone
  - d. Offers the ability to also do simple remote doctor visitations using quality video conferencing camera and audio. No issues with poor microphones or covered microphone muffled voice
  - e. Can use audio-only smart assistants in modesty rooms like bedroom and toilet. Can be a simple emergency solution for toilets to provide immediate audio linkage when required or allow fallen AIP person to call for help as augment to a pendant solution.
- iii. **Use of Cameras in the AIP person's home.** The use of Cameras in the home to allow remote visual monitoring of the AIP home is a key part of the ability for someone who is usually on their own and does not have a daily person call and spends days on their own. Cameras can be seen as too invasive for many people to have in their home watching their every move but when confronted with the choice of (i) paying for home help (ii) burdening your kids and carers to call often and spend more time checking up on you (iii) not wanting to use other sensors in the home – the camera can be the best option for the AIP person to remain at home for as long as possible.
  - a. The key to the camera solution is to use them in the context of the AIP lifecycle and in particular when people have limited movement and ambulatory problems. When AIP persons are constrained to living in their house 24x7 and do not drive or only leave the house when family members take them out they usually get constrained to spending most of their time in 3 rooms in the house. With some mobility, they spend their times typically in 3 rooms – in kitchen for meals and variation of time spent in the room, in living room typically in front of

- the TV and for change of location to kitchen, bedroom. The kitchen and Living room with camera additions give huge telemetry and contact with AIP to allow the longest periods of no visits to the AIP person or minimize the simple checkup time spent. When the AIP person's mobility is compromised they become essential to allow someone to remain in their home. They can even afford someone to spend all their time in a bed with the comfort of immediate access to their wellbeing and contact by remote carers and nursing staff.
- b. The key to the camera is to also offer
    - i. Motion and the new person in House triggers – as Carers and Kids cannot spend all day looking at Mom or Dad in their home
    - ii. Fall detection – improvements in AI is now making this possible
    - iii. Facial detection – improvements in the technology now make it capable to recognize regular visitors and carers vs new entrants to the home
    - iv. Health checking by Camera – High-resolution IR capable cameras allow the ability to detect pulse rate, potential high and low blood pressure, changes in walking gait and other medical insights
  - c. Also offer fire and other detections (my not be as good as dedicated sensors for smoke or fire)
  - d. Offer with inbuilt speakers and mic – an alternative to connect if the primary phone or smart assistant connection is not working. Many cameras now have mic and speakers built in to be able to communicate briefly into the room.

Monitoring forms the most important phase of any Aging in place solution. One key question is will 65+-year-olds or their carers see and understand the value of having or providing an AIP monitoring solution. An effort will need to be put into marketing the solution values and the importance of early adoption and not the usual ‘wait for something to happen’ before doing something about it. The other element of monitoring is to provide a simple converged solution that is packaged like a typical MSO service where it’s simple and integrated into the existing offering. The competition for an MSO monitoring solution is the Do it yourself process – so the integrated offering in the typical broadband, video, and mobile solutions of MSO will be the key driver to adoption.

The Figure 19 below shows the simple effect of the five simple MSO supplied or leveraged devices and sensory information to provide the single most important, simplest, and most effective service of AIP – the decision for carer or kids to visit their AIP parent. While visiting Mom or Dad can be a vocation for many carers and as we showed in earlier sections that 65.7% of kids are close to their parents and about 10% are very far so need remote monitoring. But as we have also shown even when carers and kids are close to parents it is a huge stress on their own lives and their time. *This monitoring solution offers as much potential to those kids as the far away remote kids – to provide as much care time as they can without having to visit because of worry of falls or other ‘absent’ information from their parent's home.* Philosophically it is a balance between their life and their parent's life and the heavy burden of guilt vs a nursing vocation drive. What this technology solution offering does is to provide the balance that seems to work with

- 100% reliable remote communication to Mom or Dad – avoid that huge frustration of not being able to contact home and worry
- Simple video triggers throughout the day – allowing an easy email or app scan to see AIP person is in good form and up and about
- Periodicity of function changes – AIP – especially when confined to house or room for much of the time gets to be very deterministic in events from the home (TV habits, motion habits, sensor value consistency) and these simple snippet updates give a fast 10-second ‘everything's ok’ view to the carer



**Figure 19 The simple monitoring flow diagram for AIP**

The second phase in the AIP timeline solution is the ‘Small Nudges and Checks’ stage where the AIP solution gets more proactive in its engagement with the home. The areas that a proactive attempt at helping the AIP process in this stage are

- Reminders and Notifications: Medicine and medication reminders, doctors’ appointments, scheduled visits to the home, bill pay reminders etc.
- Nudges and Checks: Take readings from any supplied IoMT devices like blood pressure or thermometers. Often the condition of not doing doctor visits is that the AIP person has to take their own readings continually. Automated locking of doors by adding smart locks to the doors. Leverage of Camera-based smart DoorBells, Automated turning on and off lights, Heating and Cooling use/on/off nudges and reminders, Food ordering nudges based on access to food order information, Pushed information from their carers and children – display on TV

Solutions already exist for these types of notifications and reminders but are usually SmartPhone driven and in recent times there have been extensions to smart assistants to also set up important reminders. The issue is a typical one that is also found in IoT generally in that it is

- Not fully understood by the population and demographic of 65+ AIP person and often not even by their Kids or Carers that these tools exist
- Still relatively complex for the majority of people
- Disaggregated and not consolidated in a simple Consumer Experience

For notifications and reminders, the Cable Operator has a better solution potential than the SmartPhone applications and that is the leverage of the TV as the main portal interface for all AIP communications. The ability to be able to access this large-screen format to affect these Aging in Place applications removes much of the inertia of the 3 issues above.


- Cable Operator provided a portal solution that can mix in AIP notifications, Video Conferencing in the same CX solution for the lean-back video experience
  - o Can leverage the directions of the 2 most popular middleware solutions on STB/SMD to add services and applications
    - AndroidTV and App Store and other AndroidTV services

- RDK V and the evolving improvements in the Downloadable Application Container solutions for both STB/SMD and Broadband devices
- Made simpler to work in lean-back mode utilizing simplest input devices
  - Aging in place adapted remote controls for ease of use. In some cases, it can be a large button remote for hard of sight elderlies
  - Use of Far-Field voice to provide input to the services and replies
  - Larger Font displays through the accessibility feature available in RDK and AndroidTV for sight-impaired AIP persons
  - Higher Audio playback for hard of hearing AIP Persons
- Even the control of the AIP persons STB/SMD by the remote Carer or family member to get the right channel or program is a well-traveled issue that many have seen when they are caring remotely for their parent(s)

Perhaps the most important nudge of all is the reminder for taking medication which in itself the feature dwarfs all others as the largest payback for the AIP process. Many AIP persons forget their medication or take it incorrectly. One simple scheme is to couple the notification of medication schedule with the TV experience and uses it to prompt for medication or taking readings from supplied IoMT devices. The TV screen and its engagement with AIP persons is a key opportunity to be able to provide all AIP services on the 65” palette that is the TV.

The last simple role the Cable Operator can play in the Aging in Place solution is the ability to help in the entrance and tracking of people into the AIP home. This is one of the key issues of the AIP process especially when there is a lack of mobility and hearing and the process of answering the door is both a difficult process and also opens security issues. Even when using solutions like food delivery when the AIP person is no longer driving on leaving the home un-aided the ability to answer the door can be the main inertia to the person staying on their own. So, a simple addition of a smart front door lock and camera solution tied to the Cable Operator solution and allowing

- Visibility of who is at the door – not only to AIP person but also to their carer or kids
  - On the TV screen and the smartphone app for both the AIP person and Carer/Kids
- Ability to open the door from the TV remote, voice or smartphone application – both for the AIP person themselves and the remote Carer/Kids
- Logging of open and close door events including thorough checking of the locking process of the door after people have left. Potential separate camera verification.



**Figure 20 The 5 key tenets of simple AIP solution set for Cable Operators**

With these 5 simple features, as shown in Figure 20, integrated into the Cable Operators current offerings and devices, there is a high probability of covering the key features that are required to provide the tools to the AIP person(s) to be able to control their own lives, reduce the burden on their Carer and Family circle and provide the realtime transparency to the home so that any event that happens that threatens their health is responded to immediately – either through direct voice and video access to the home or the dispatch of someone close to the home. This solution works really well for local carers and family members to keep their time optimized for f2f visits and removes much of the guilt of family members who are not close to their AIP parents by being able to cost-effectively direct local resources to help only when required or in a normal deterministic schedule.

The cost of these solution elements is also very low from both a capital investment perspective for both the Cable Operator and the Aging in Place home.

For the Cable Operator, the cost elements overlap a lot with the capital investment for the broadband and video and phone experience and the main constituents are

- Broadband Gateway with Wi-Fi – the **amortized cost for Broadband Service sale**
- Smart Media Device – **amortized cost with Video entertainment experience**
- Smart Phone Device – **amortized cost with Mobile service if offered by Cable Operator**
  - o All AIP App elements can be downloaded on a third-party device
  - o AIP person and their Kids/Carers also download and use Apps as part of service
- Operator supplied
  - o Camera(s)– Prices can range from \$25 to \$900 depending on solution quality
    - Camera additions to Smart Media Device for lean back Video conferencing
    - Camera additions to other non-SMD locations
  - o Additional in-room audio smart assistants – Prices can range from \$55 to \$180 depending on the configuration
  - o Smart Door Bell including camera – prices can range from \$120 to \$300 depending on configurations
  - o Smart Lock for Door – prices can range from \$150 to \$300
  - o Optional LTE/5G backup devices for Broadband services
- Cloud storage and transaction costs for device telemetry and cloud to cloud partner and Care Portal engagements
- Additional Technician costs for installation of AIP additional devices in particular Front Door solutions of Doorbell and Camera
  - o There is the scope that these can be done as Self Install Kits given the reasonable simplicity of Doorbell and Lock installs
  - o Internal Camera installs can all be self-install for AIP or Family members to do


## 4. The Cable Operators solution play for Aging in Place

This section along with Appendix A, will define the recommended AIP Home architecture and the role the Cable Operator can play to provide a valuable service to this market where the price of staying at home as long as possible can sometimes be not even defined in monetary terms. Finding the balance of

- (i) Existing CAPEX spend and OPEX spend on install and support for Broadband Video and Mobile services
- (ii) Adding in additional elements to the existing Broadband Video and Mobile devices to support the AIP market and to abstract this for maximizing the almost 9% of your network

- and 55M people in the US that can avail of the service. Sharing this CAPEX investment is key to the business proposition
- (iii) Looking at the extension of Aging in Place investment and it's additional overlap to the Telemedicine business for **EVERY** customer on your network. Aging in Place solution and Telemedicine services have a lot of overlap with investment and partnership opportunities
  - (iv) What the Operator provides vs Care Companies, Insurance Companies, Health Providers and Government support agencies for improving aging life and reducing the cost of out of home care
  - (v) Providing client devices that add to the support of Aging in place but minimize the medical nature of the offered AIP service. Leveraging a Smart Home and Security solution to also support key non-medical Aging in place services
  - (vi) Leveraging existing Cloud services to create the Cloud to Cloud telemetry connections to specialist Aging in Place partner solutions.
  - (vii) Leveraging the change in services architecture paradigm of moving away from locked Operator services on Broadband and Video devices to one where containerized service can be added to the connectivity network to effect new services through the devices and endpoints in the home.

Figure 21 illustrates the key eleven elements of the Cable Operators solution for *AIP Home*. These solutions focus on the simple non-medical elements that make Aging in Place work with the virtual access to the home by carers and family. It also allows the Cable Operator to partner with AIP partner companies offering resources and services to the AIP process ranging from - home help services, food and meal preparation services, medical services, insurance services, specialist care services, etc.




**Figure 21 Eleven key service offerings from a Cable Operator to support AIP Home**

A detailed discussion on these eleven services, their needs, the opportunity, the cost (CapEx and OpEx) implications, and the application extensions (such as telemetry) are provided in the Appendix.

## 5. Business case with estimates of tiered service prices

Before putting together the business case for the AIP solution proposition for a Cable Operator, let us look at some assumptions –

- The average life expectancy in the US is 78.6<sup>13</sup> years in 2020 and it is trending up: The likelihood



**Figure 22 Probability of dying between X and X+1 years**

of people dying or leaving their home is part and parcel of the AIP business case process. When you are selling solutions and investing CAPEX in the 65+ customer base you have a much higher probability of any invested capital or install costs for recurring revenue that may not reach their full potential. The life expectancy of 78.6<sup>13</sup> for the US population determines your strategy. Figure 22 the probability of dying at different ages. This gives an idea of how the aging population and their supporting team think about different investments to the *AIP Home*. So, for risk mitigation, it probably drives a balance of getting some upfront capital investment costs for any solutions offered and a recurring revenue opportunity.

- The financial stability for the aging population from a Lifetime Income<sup>14</sup> point of view is essential for the future *AIP Home*: Increasing life expectancy and lack of proper financial planning

will severely impact the lifestyle of the aging population. Supporting such financial services and also bundling the technical solutions to meet with the *AIP Home* is essential for the wide-scale adoption of the services that we discussed in this paper.

- The fragility of the aging population forces innovative and expeditious services: Injury like a broken hip (300,000 seniors break a hip in the US each year of which 70% of them are women) forces the aging population, especially *AIP Homes*, to adapt to their conditions. Studies have shown<sup>15</sup> that the mortality rate of a broken hip over 65 years of age doubles over 12 years after the broken hip. In usual care, the reported 1-year mortality after sustaining a hip fracture has been estimated to be 14% to 58%. The relative risk of mortality in the elderly patient population increases by 4% per year. The first year after a hip fracture appears to be the most critical time.

The above elements reinforce that AIP investment should be either (i) some upfront cost for CAPEX and monthly lease fee for equipment and service (ii) buy out option of the equipment with a monthly fee for service. The affordability of the technology assists in remaining at home is somewhere between

- Pay anything to remain at home and reduce the burden on kids and remain independent
- The average income of the 65+ age group who are not working is \$1,700 per month
- The average cost of Home Health Care of \$4,000 per month deducting what Medicaid offers to support for Home Care<sup>16</sup> such as, home and environmental accessibility modifications (alterations such as wheelchair ramps, walk-in bathtubs, stair-lifts, and environmental aids for lighting), medical equipment, and supplies, Personal Emergency Response Services (PERS) are electronic monitoring or call and respond services that enable persons to live alone or to spend portions of their day without direct supervision. There are four different categories of programs within Medicaid that offer funding that can be used to pay for electronic safety monitoring for the

<sup>13</sup>NVSS, *National Vital Statistics Reports*, June 2019, available [here](#)

<sup>14</sup> American Academy of Actuaries, *Risky business: Living longer without income for life*, June 2013, available [here](#)

<sup>15</sup> Geriatric Orthopedic Surgery and Rehabilitation, *The 1-Year Mortality of Patients Treated in a Hip Fracture Program for Elders*, Sept 2010, available [here](#)

<sup>16</sup> Paying for senior care, *Medicaid's home care benefits: Eligibility, waivers and application information*, August 2020, available [here](#)



elderly (therefore for PERS / medical alert services) (Medicaid Waivers, Consumer Directed Services, Medicaid State Plan – Personal Care Attendant, Money Follows the Person. There are also non-Medicaid sources of financial assistance for PERS devices<sup>17</sup>. Some of them are provided in the table below.

- There is also a high probability that Family Members will also help out with the costs for technology-based AIP solution. They have the greatest to gain by minimizing their face to face visits to their parents to when they want to go versus having to constantly check if their AIP parent is ok. They may also be able to continue to claim the carers allowance themselves while using the technology packs to make their visit times much more efficient.

| Service Type  | Description                                                                                                                                  | Cost                                            |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------|
| Basic Service | Wearable pendant to a call center for emergency response                                                                                     | \$25 - \$50 / month                             |
| Mid Range     | Pendant / watch with automated fall detection, call center, emergency response, family notification                                          | \$30 - \$60 / month                             |
| High End      | Wearable device, multiple in-home sensors, two-way communication, online reporting, emergency response, multi-party notification and add-ons | \$500 - \$1000 startup and \$50 - \$100 / month |

The current costs for these PERS services are defined above. The AIP solution being proposed in this paper does cover the PERS safety monitoring requirement and extends beyond this with other functions that really cover the stay at home independence inertia items like Door Ingress and Egress. The ability for both the camera and the drop in audio and video services as well as additional features in the remote

### A note on AIP technology adoption

Age related issues *don't really respect the socio-economic demographic*. While more money in retirement can help you pay to stay healthier – a serious fall levels everyone. Other degenerative diseases also drive same decline no matter how much money you have.

For technology based assist solutions there is a belief that the *adoption* of this type of service will start in families that have that *bit more money and are already tech-savvy* having seen the benefit of consumer IoT devices and smart assistants. In fact many of these families may try and help their AIP parent(s) with Do it Yourself Technology solutions. We would like to think however, that these proposed AIP technology solutions can find their way to the poorest demographics as they may have the largest impact in keeping these seniors in the own home and out of the stressful work of assisted living, day care and other hugely expensive burdens.

Its also hoped that *the Medicare and Medicaid services can also look to these technology solutions* as driving large savings in their premium payouts and extend PERS and other coverage items to include the suggested device sets indicated in this paper. Why is a DOCSIS GW , WiFi AP , Extender or IoT hub not a potential insurance subsidized device given roles in preventing injury and providing reliable immediate triggers when threshold of sensor readings are breached and doors are opened. More work needs to be done to move more of the technology assists in the home for AIP into subsidies from insurance companies.

control unit – all improve the ability to do PERS.

<sup>17</sup> Paying for senior care, *Medical Alerts & Personal Emergency Response Services Costs*, May 2020, available [here](#)



We believe, as highlighted in the insert “*A note on AIP technology adoption*,” this market opportunity will take off for the Cable operators. Let us look at the business case for the Cable Operator in Aging in Place. For the purposes of this paper, to make it simple, let's assume

- Margins on the proposed services below are even better than Broadband margins
- Because of the lower survivability rate in this age group (65 and above), we have to reduce risk on Capex and Opex investment
  - o So, most models will be upfront installation cost with recurring lease cost on device
  - o The AIP services as a Software Service will be licensed as yearly maintenance costs for the additional software updates and improvements
  - o Some higher initial onboarding costs for third party OTT AIP service providers – the assumption being that a limited number of partners per region or state will be used
- The services below do not include the cost of basic triple-play services and the device lease costs.

So the following analysis outlines 5 simple illustrative AIP Packs. The potential revenue opportunities are highlighted in the table below.

| AIP Service Offering Costs                                                         | Install cost | Upfront cost | Leasing cost/mo. | Purchase cost |
|------------------------------------------------------------------------------------|--------------|--------------|------------------|---------------|
| <b>(1) Family communications pack</b>                                              |              |              |                  |               |
| SMD + Video conference                                                             | \$50         | \$200        | \$10             | \$500         |
| Additional SMD + VC                                                                | \$50         | \$200        | \$10             | \$500         |
| <b>(2) Whole home family communication adder</b>                                   |              |              |                  |               |
| In room smart audio assistant + IoT                                                |              |              | \$8              | \$99          |
| Two pack - in room smart audio assistant + IoT                                     |              |              | \$14             | \$180         |
| <b>(3) Entrance and egress security kit</b>                                        |              |              |                  |               |
| Smart doorbell                                                                     | \$60         |              |                  | \$150         |
| Smart doorlock                                                                     | \$120        |              |                  | \$250         |
| Smart doorbell + doorlock pack                                                     | \$150        |              |                  | \$325         |
| <b>(4) Camera</b>                                                                  |              |              |                  |               |
| 1080P only single camera                                                           | \$50         |              | \$7              | \$120         |
| Two pack 1080P only cameras                                                        | \$75         |              | \$12             | \$225         |
| 4K single camera                                                                   | \$50         |              | \$12             | \$200         |
| Two pack 4K cameras                                                                | \$75         |              | \$20             | \$390         |
| <b>(5) AIP for IoT pack</b>                                                        |              |              |                  |               |
| Motion sensors                                                                     | \$100        | \$100        | \$4              |               |
| Smoke and fire detection                                                           |              |              |                  |               |
| Kitchen sensors                                                                    |              |              |                  |               |
| Toilet flush sensor                                                                |              |              |                  |               |
| <b>(6) Skills support</b>                                                          |              |              |                  |               |
| (1) + linkage to TV and TV notification                                            |              |              | \$15             |               |
| (2) + linkage to TV, remote, voice system                                          |              |              | \$20             |               |
| (5) solution software                                                              |              |              | \$20             |               |
| AIP care portal connect service (Care portal to add services to devices installed) |              | \$250        | \$50             |               |

#### (1) Family Communications Pack:

This service focuses on a simple large-screen video conferencing integrated into the TV experience. Features include - Smart Media Device and Camera bundle, Smart Phone App for Family and Carer Givers, Features like TV pause when a call, Family members can leave notifications on the TV screen, etc. This package can be for one room or more than one with additional service and equipment

lease. It offers options for complete buy out of the device vs lease cost.

- (2) Whole-Home Family Communications Adder: This service adds Audio only Smart Assistant adders to additional ‘modesty’ rooms.
- (3) Entrance and Egress security kit: This service includes, Smart Doorbell and Smart Lock combination for TV controlled secure visitor ingress and egress to the home. Note that the installation costs are higher in particular for the Door lock.
- (4) Cameras: Add 1080p and 4k camera options added to the SMD camera solutions. These can be placed in other rooms or on outside locations for TV-based security assessment
- (5) AIP IoT pack (for non-medical devices): This service focuses on non-medical related IoT packs on giving the Family and Carer simple daily life functions transparency. These devices include - room motion sensors, sleep sensors, toilet flush, kitchen sensors (fridge, microwave, and hob/oven usage), in room location, smoke heat gas detection, etc.

All of these are tied to the Family and Carer Giver app and are displayed on a TV. These skill packs are charged additionally to the device costs/leasing costs and will have constant innovation on them as well as requiring cloud resources that will cost on an annual basis.

| Sample pack contents                    | Upfront cost | Purchase cost | Per Month | Take rate | PM Revenue | PY Revenue |
|-----------------------------------------|--------------|---------------|-----------|-----------|------------|------------|
| (1)                                     | \$200        |               | \$25      | 10%       | \$3        | \$30       |
| (1) + (3)                               | \$200        | \$675         | \$45      | 20%       | \$9        | \$108      |
| (1) + (3) + Two pack 1080P (4) + (5)    | \$300        | \$825         | \$81      | 25%       | \$20       | \$243      |
| (1) + (2) + (3) + Two pack 4K (4) + (5) | \$300        | \$925         | \$103     | 30%       | \$31       | \$371      |
| Above + AIP Care portal                 | \$250        |               | \$133     | 15%       | \$20       | \$239      |

\$83 \$991

How will this then roll out to the AIP population with different requirements depends on privacy, different budget levels, etc. The above simple table outlines 5 configurations of the above packs from a simple one SMD VC communications pack to a premium most of the services selected in more than one room pack. These range in price per month from \$25 to \$133 per month. A Take rate estimate of the 5 packs is shown in the above table with Pack 4 is the most popular at a 30% take rate. Using this trivial simple take rate – the average cost per month is about \$83 per AIP home. Note there is also upfront payments additional to this per month price as well in our proposed model.

This is in line with research data and PERS costs shown above (this solution goes beyond basic PERS features) that show AIP persons or their family willing to pay at \$50-\$99 per month to remain as independent as possible at home.

| Target market (yearly revenue projections)   | Market size | Take rate | Revenue (Billions) |
|----------------------------------------------|-------------|-----------|--------------------|
| AIP Homes (Total market for 65+)             | 56,000,000  | 3%        | \$1.67             |
|                                              |             | 30%       | \$16.65            |
|                                              |             | 50%       | \$27.75            |
| Elderly on their own need help               | 3,750,000   | 3%        | \$1.11             |
|                                              |             | 30%       | \$1.12             |
|                                              |             | 50%       | \$1.86             |
| > 65 households in 2M subs                   | 170,000     | 3%        | \$0.01             |
|                                              |             | 30%       | \$0.05             |
|                                              |             | 50%       | \$0.08             |
| Elderly on their own needing help in 2M subs | 45,000      | 3%        | \$0.00             |
|                                              |             | 30%       | \$0.01             |
|                                              |             | 50%       | \$0.02             |

As a final step to give a quick feel for the overall opportunity for revenue let's look at a very simple model of 3%, 30% and 50% take rates of the 2021 AIP population across

- The entire 56M AIP Homes base
- The smaller subset of Elderly on their own already needing ambulatory and other help – 3.75M
- And a simple hypothetical 2M subscriber MSO for their revenue return

As you can see below it could be a \$28Bn per year market if 50% of homes over 65 had such a Cable Operator supplied solution. Even with those needing help now at 50% penetration, it's a \$1.9Bn per year revenue opportunity. For a 2M subscriber Operator with 50% of their 65yr+ homes taking a solution – it's an \$84M per year opportunity. And if only 50% of the AIP people already requiring help take it – it is a \$23M per year opportunity.

Of course, there are other elements that can be further considered, such as -

- Additional Broadband and Video services for new subs or upgrades to existing ones
- Stickiness – AIP person is likely to never leave Operator as Broadband and Video Provider

- Stickiness – their family and carers may also switch to the Broadband and Video services of their AIP parents if they see the additional value. And there could be family packs offered at some overall discount as a new tier of service
- Additionally, AIP services that go beyond that are only technological specific

## 6. Conclusions

As Cable Operators enter the next decade between 2020 and 2030 there are a number of key trends correlating to allow them to add new high-value revenue opportunities. As Operators move to the ‘quad-play’ of Voice Video Data and Mobile – there is another level of Service that is emerging in the Aging in Place and Telemedicine space. This area offers high value per month returns to the Cable Operator and potentially \$82 per month revenue from 9%+ of its subscriber base. With the advancement in IoT devices, Far-Field Voice and Smart Assistants, and in particular the evolution of the STB to the Smart Media Device a Cable Operator has a device arsenal that can be re-purposed and in some cases double-dipped to Aging in Place revenue opportunity. By choosing a simple path based on connectivity, SMD, home communication, home ingress, and egress, and home telemetry solution and extending to partner with specialist AIP companies the Cable Operator can provide a longer time at home for many of their Elderly and aging customers. Not only is it a profitable solution direction but a hugely worthwhile solution space that can improve the lives of up to 55M people in the US but also bed in a new digital home solution that will also grow to improve Telemedicine and other home services.

# Appendix A: Some solutions operators can offer for AIP Homes

| Service offering                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Opportunity information                                                                                                                                                                                                                                                      | Additional AIP needs                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <u>Broadband GW and Wi-Fi</u> : Fundamental components that can be primary internet services to 20% of the <i>AIP Homes</i> that do not have an internet connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 20% of the AIP Homes<br><br>CapEx is absorbed in operator spend.<br><br>No additional OpEx.                                                                                                                                                                                  | Option to use the connectivity telemetry and data from the home and apply it to AIP trend analysis. Simple ML-based trending of internet usage (no use – alert carers), health and status of connected devices, changes in numbers of connected devices, new devices present in the home. |
| <u>LTE/5G backup</u> : The solution to resolve an outage on the HFC network is to supply a backup LTE/5G solution that can switch over reasonably fast and support the WAN internet connectivity for all the in-home AIP services and communication. This can be charged on a per-use basis to minimize the cost of LTE and also to make it easier for those AIP families to purchase the service when worried about the cable network dropping.                                                                                                                                                                                                    | For those families looking for beyond 5x9's reliability of connectivity to an AIP place home.<br><br>No additional CapEx<br><br>Requires OpEx through additional training and installation from a technician (could be self-install) and support costs and training.         | Needs additional support for the health and provisioning of the LTE device as well as a solution to be able to reliably swap access support from Cable to LTE and back again. No real specific AIP elements are required.                                                                 |
| <u>SMD and TV as the main AIP portal</u> : The STB is on a new evolution trajectory to move from being just a traditional decoder to being an in-room device that can open up the 65" TV as the portal for many services including AIP. The device itself and the location of TV's in the AIP home – correlates with the two main rooms that the AIP person spends most of their time outside of the bedroom. Typically these rooms are (i) kitchen and (ii) living room – which as mobility decreases in AIP lifetime the AIP person spends more time sitting and in front of the TV. This makes the SMD the idea device to become the AIP command | As the Cable Operator moves from humble STB to the Smart Media Device solution and brings in new Smart Media devices – they can enable an AIP service to the device easily with the addition only of a Video Conferencing camera as an additional service over their regular | The opportunity to be able to leverage the SMD Portal to both the TV pixels and speaker audio to be able to connect AIP resources to the AIP person.                                                                                                                                      |

| Service offering                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Opportunity information                                                                                                                                                                                                                                                                                       | Additional AIP needs                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| console for both the AIP person and their carers to connect and communicate to minimize the burden of face to face visits and care. With the use of Videoconferencing, it also makes it much easier to have immediate and frustration-free discussions with parents to improve their psychological well-being.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <p>next-generation TV experience evolution.</p> <p>Potentially Cable Operator can absorb Smart Media devices in its evolution to IP Video over Wi-Fi CapEx and adds more Far Fied voice and visual smart assistant support.</p> <p>Requires additional training and support for the SMD and AIP services.</p> |                                                                                                                                                                      |
| <p><u>Simple remote control:</u> As we age our ability to use and work with devices that can be lost or misplaced, are hard to see, have small buttons, have complex button presses, etc. These can be a source of massive frustration as well as potentially stop the ability to accept incoming calls or simply change the channel to watch a favorite program. Changing the design of the remote control to be <i>AIP Home</i> friendly such as be less easy to lose or slip down the side of the chair, have bigger and fewer buttons, support push to talk voice to enable services and smart functions, have a remote from home initiated ‘find me’ beep or noise function to locate it.</p> <p>The remote can also be repurposed to be a pendant remote control that can also provide emergency alert services. Idea’s for this service include: Adding accelerometer and <b>gyrometer</b> support to the remote to detect falls and rapid movements – with an automated call out to Carer or Family member, Red button press to engage carers or specialized call out services, etc.</p> | <p>Opportunity to create a differentiated – fit for function – AIP friendly remote control device.</p> <p>This would be a separate SKU of remote for AIP purposes.</p> <p>Requires additional training and support costs.</p>                                                                                 | Needs additional support for the services added to the remote – particularly if its IP addressable and supports the ‘find me’ and emergency alert services.          |
| <p><u>In-room speaker/intercom:</u> To complete the coverage of the <i>AIP Home</i> transparency and visibility for Carers it is important to add additional communication devices to the ‘modesty’ rooms in the home like a bedroom and possibly toilet. This completes the formula of being able to constantly contact the person and also afford them the potential to constantly contact their carers or family as they move around rooms. The addition of an audio-only smart assistant will complete the</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <p>The opportunity here is to introduce general whole-home voice and audio for the general population but to also drive this for an AIP specific application. The operator supplied Smart Assistant solutions</p>                                                                                             | Needs additional support for the IoT and Voice/Audio data models to support management and provisioning of the device. AIP applications can be applied to the device |

| Service offering                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Opportunity information                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Additional AIP needs                                                                                                                                                                                                                                                                                                                             |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>transparency coverage of the home and make the <i>AIP Home</i> very much a virtual extension of the carers and families home – with a minimum burden and maximum return.</p> <p>The Operator can also affect another level of security and privacy over these audio listening devices and also include easier point to point access to designated only listeners or call endpoints – typically just the AIP carers and family members.</p> <p>The device can also offer additional IoT radios to also engage with non-BLE solutions added to the home for smart IoT or even non-BLE IoMT solutions.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <p>have additional advantages for home use.</p> <p>Absorbed in normal Cable Operator Capex spend : Potentially. The addition of Far-Field Smart Speaker based solutions as an augment to the SMD Visual Smart assistant is an architecture that has merit for creating a new whole-home platform for new services including Aging in Place services.</p> <p>Requires additional training and installation from technicians (could be self-install) and support costs and training.</p> | <p>– but typically will be cloud hosted and leverage skills frameworks that have already been integrated for regular consumer services.</p>                                                                                                                                                                                                      |
| <p><u>Camera solutions for room visibility in defined areas:</u><br/>Cameras in the home cause lots of discussion about invasion of personal privacy. However when confronted with having to -leave the home for Assisted living or Nursing home, add more burden to family and carers with more visits to <i>AIP home</i>, increased cost in home health care, worry about being on one’s own so outdoor camera solutions can give peace of mind, these concerns will lessen. The addition of Cameras to the <i>AIP home</i> – is a much more palatable and acceptable solution to the AIP person when it adds this value.</p> <p>Cameras in the home for AIP should be allocated in the following ‘modesty’ fashion - they should only be added to rooms like Kitchen and Living Room , they can be combined in those rooms as both constantly viewing devices and video conference solutions – especially when combined with SMD and TV, they have additional role to provide a view of front door to the TV and Smart Phone for automated visitor access to the home, their telemetry can be also used for other AI based and AIP and Telemedicine services using video to extract meaningful home insights</p> | <p>New class of device and service that has multiple roles across the general customer services but a specific relevance to Aging in Place services.</p> <p>New chargeable device as part of AIP or Security initiatives. Additionally, options for charging for recording or AI services on the data as part of menu of services on video telemetry.</p> <p>Requires additional training and install from technician (could be self install) and support costs and training.</p>      | <p>Video and Audio recording services as well as events like motion , person and other finer grained AI derived services. AIP applications can be built on the Video services and solutions that can analyze heart rates, changes in walking gait, falling as well as fire , smoke or other solutions can be overlaid on the data collected.</p> |

| Service offering                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Opportunity information                                                                                                                                                                                                                                                                                                                                                                                                 | Additional AIP needs                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>For the Carer and Family cameras - provide more value to them than the AIP person and so are their main vehicle to establish wellbeing of AIP person, provide a fast way to track the activities of the day like eating and drinking and visual verification of taking medicine, provide a way to approve ingress to the home for visitors remotely with confidence – a scheduled visit or repair service etc., provide emailed or Instant Message event triggers with captured frame or video snippet for activity in the room – making it easy to verify the key events of the day of (Up and going, in the relevant rooms, moving, time based changes of rooms, visitors etc.)</p> <p>The Operator can supply these cameras solution as part of Video Conferencing service on SMD as well as extra cameras for Door and other rooms as part of their service. Camera solutions scale outside the AIP opportunity to other home peace of mind security services and in particular of infrequently occupied second home monitoring services. The device can also offer additional microphones and speakers to also support camera based intercom features. The Camera should also support IR and Nightvision modes to be able to support visibility through 24 hour period. Cameras will also trigger alerts at any fire's starting in the home as motion events.</p> |                                                                                                                                                                                                                                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                |
| <p><u>Camera doorbell</u>: One of the limiting factors of Aging in place that is solvable with technology is automated entrance to the home for visitors and also answering the door to strangers when desired. This is a particularly important feature when mobility has reduced of a single Aging in Place person. Allowing delivery of food with trusted people as well as scheduled maintenance or health workers using a secure automated door solution is a key part of elongating the stay at home battle of aging in place. Adding a Camera doorbell provides this capability and can uniquely be tied to the STB and SMD by the Cable Operator to allow someone who spends their time in a chair or bed to also control entrance to the home.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <p>Opportunity : New class of device and service that has multiple roles across the general customer services but a specific relevance to Aging in Place services.</p> <p>Absorbed in normal Cable Operator Capex spend : New chargeable device as part of AIP or Security initiatives. Additionally, options for charging for recording or AI services on the data as part of menu of services on video telemetry.</p> | <p>Video and Audio recording services as well as increasing ability to do facial recognition on the door bell camera. Logs of entrance and egress to the home as well as being able to deal with unwelcome people at the front door. Gives also an additional level of security as people don't loiter when they know they are on cameras.</p> |

| Service offering                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Opportunity information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Additional AIP needs                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><u>Smart door lock:</u> Works in conjunction with the Smart Doorbell to provide a good secure automated solution to grant secure access to approved visitors. With key elements of aging in place being - food delivery, health worker visits, ambulance visits to take you to scheduled doctor visits, amazon or online services delivery etc.</p> <p>The ability to automatically open the door by the AIP person or even remotely by the Carer and Family can be a huge help in remaining independent in AIP timeline.</p> <p>There are additional devices that can be fitted to the door to make sure it closes when opened as well as the door lock telemetry for close to ensure highest security levels on the door closure process.</p>                                                                                                                                                                                                                                                                                                                       | <p>Additional Opex costs :<br/>Yes – requires additional training and install from technician (could be self-install) and support costs and training.</p> <p>New class of device and service that has multiple roles across the general customer services but a specific relevance to Aging in Place services.</p> <p>New chargeable device as part of AIP or Security initiatives..</p> <p>Requires additional training and install from technician (could be self-install) and support costs and training.</p> | <p>For security purposes the application of software solutions to ensure the door is properly closed and the door/open close correlated to the doorbell telemetry ensures complete visibility to the door access solutions.</p>                                                                                                                                          |
| <p><u>Partner supplied IoMT pack:</u> There is a place for the Cable Operator in the medical devices that can be required in the Aging in Place journeys. Typically, they will be supplied by specialized Health Care Companies, Insurance companies, Specialist Aging in Place solution providers and in many cases come pre-provisioned for immediate send of information to a care portal database. The Cable Operator can add value to this service by - providing the IoT hub elements to connect the IoMT devices (Typically BLE but also offering Wi-Fi and ZigBee and ZWave onboarding, doing BLE to IP conversion from the device, supporting reliable first time pairing of the devices to the IoT hub and connection to the cloud etc.), displaying the status and onboarding and results of the devices integrated into the both (The SMD visual display on the TV, any Operator supplied Smart App, integrating the IoMT devices into the Cable Operators TV Consumer experience, Adding audio playback of the recorded values for sight challenged AIP</p> | <p>Partnership opportunity with specialist monitoring and Aging in Place care companies. Cable Operator provides the Hub function to the IoMT device pack issued by the partner monitoring company. Cable Operator takes a fee for the onboarding service and integration into SMD visual path for notifications and other services.</p> <p>An OTT service that can be added to the Cable Operators service architecture for other services like their own IoT</p>                                               | <p>The solution should probably be offered as Virtual Network Function added to the Cable Operators network and together with containerized application solutions added to the Broadband and Video Service delivery system – offer the Cable Operator the option of being able to partner with specialist Caring Services and Resources. It is not expected that the</p> |



| Service offering                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Opportunity information                                                                                                                                                                                                                                                                                                                                                    | Additional AIP needs                                                                                                                                                                                                   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>person(s)), adding additional security and redundancy via the addition of LTE/5G WAN backup.</p> <p>Opening up cloud to cloud interfaces between the Connectivity path and the Medical Cloud to provide secure reliable integration with the home User experience. This is typically done by - Cloud to Cloud interface definitions, containerized applications that can be added to the SMD and IoT hubs to gather IoMT information, providing logging solutions to ensure trace back of readings and frequency.</p> | <p>services and even Wi-Fi management. Additional effort required if Cable Operator provides more features in the connection of IoMT to the specialist Aging in Place providers cloud.</p> <p>No additional or at least very minimal as part of the next phase of addition of containerized services into the Cable Operators Broadband and Video and Smart Solutions.</p> | <p>Cable Operator does anything with any data from the IoMT devices except to provide tools to ensure they have onboarded and remain connected as well as other alerts for any outages or unscheduled disconnects.</p> |

Cable operator cloud: As mentioned in earlier sections – the expected solution for Cable Operators AIP business is to provide additional cloud support to the (1) Cameras and Smart Assistant based services – offering additional access to these devices directly to AIP person(s) and carers and family. There can be opportunities to also allow specialist AIP companies to access cameras with permissions from family and AIP person(s). These services can be provided cloud to cloud (2) Doorbell and Smart Lock solutions as extensions of any Smart Home or Security solutions offered by the Cable Operator. The Cable Operator providing the automated entrance service to the AIP person is a cloud hosted application and needs additional services to ensure closure and alerts to AIP family and carers if closure contact is not made. The specific IoMT pack is not typically controlled by the Cable Operator and is typically a pass through service that the Operator will offer to the AIP either through partnered solutions or potentially to designated solution offerings recommended by the family or AIP themselves. The definition of the interfaces to these third party IoMT monitoring and care portal companies has no specific standard at the moment but there are several potential solution architectures for this that could provide some future telemetry and information sharing from Cable Operator domains to AIP resource domains. For now the simple mechanism is to send all the IoMT information through the Cable Network to the defined cloud endpoint after VNF and

| Service offering                                                                                                                                                                                                                                                                                 | Opportunity information                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Additional AIP needs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Containerized applications for AIP Resources are installed in the Cable Operators Service Delivery Platform.</p>                                                                                                                                                                              | <p><u>AIP partner clouds and cable operator interfaces:</u> This area demarcates the non-Medical and Medical elements of Aging in Place. In the paper we have outlined the value areas that the Cable Operator can add from its own cloud to home AIP services. However there is another important opportunity that can be also added to the Cable Operators services to the AIP home. The specialist Care Companies, Medical Resources and other Aging in Place resources can connect either - Directly through the operators network to the devices, Support some potential for Cloud to Cloud integration for Operator supplied value add services to the medical or other AIP OTT services</p> | <p>This offers the potential to be able to – (1) Onboard with Operator supplied Hub devices (for its own AIP and general Home IoT and Smart Home services) rather than additional hub. New Software Delivery schemes unlock the possibility for a containerized AIP medical app to be added to support Aging in Place medical device packs and other internet connected services above the basic ones supplied by the operator. (2) Offer additional Operator pulled telemetry and even Machine Learning to the AIP resource cloud service (potentially driving some standardization of data sharing) and combine this information for a better overall service from the specialist AIP service provider</p> |
| <p>This is a more complex proposition for the Cable Operator as there are many AIP specialist companies, and they differ regionally throughout the US. In most cases the AIP resource company is recommended by - Doctor/Hospital, Rehab center, Day Care center, Health Insurance Provider.</p> | <p>Or researched by Family member or AIP for the Home Help or monitoring solutions.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <p>But the potential does exist to be able to onboard and integrate with these solutions typically via - software architectures to allow downloadable services in Operator supplied devices, cloud to Cloud integration</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Service offering                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Opportunity information | Additional AIP needs |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|----------------------|
| <p>and definition of interfaces to share data and information.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                         |                      |
| <p>More research and ‘analysis needs to be done in this area of interfacing with Aging in Place medical and other specialist services.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                         |                      |
| <p>Data exchange and the privacy of data will also need to be expanded and is outside the scope of this paper. While there is the HIPAA and PHI specifications to govern access to sensitive medical data generally all data from the AIP person and home will start to become more privacy focused – particularly when simple sensor information combined with medical devices can be merged in an overall view of AIP persons health levels. There has already been a new act created in California – the California Consumer Privacy Act – CCPA which is similar in many ways to the trend in Europe with the General Data Protection Regulation being put in place to define rules around consumer privacy in a digital world.</p> |                         |                      |
| <p>Before going onto the next section we will also mention that there are many other additional sensors that could be added to ensure more safety and peace of mind visibility in the home. These include but are not limited to - motion sensors, smoke and fire detection, sleep sensors, toilet flush sensors, fridge open sensors, door open sensors etc.</p>                                                                                                                                                                                                                                                                                                                                                                      |                         |                      |
| <p>These can be added as required to the Cable Operator Aging in Place pack and the telemetry added into the flow for the Carers and Family members to interpret. These devices can provide safety and security features as well as general indications of healthy hygiene practices for the AIP person. Declines in the normal levels of feeding, toilet frequency and sleep can be used to pre-empt medical issues and get earlier family and doctor interventions. These statistics can also be combined with supplied medical devices telemetry to get a clear picture of overall health and aid prognosis by Doctors.</p>                                                                                                         |                         |                      |

# **Why Are Cable Operators A Natural Fit To Support Telehealth**

## **An Inter-Industry Perspective**

A Technical Paper prepared for SCTE•ISBE by

**Dr. Sudheer Dharanikota**  
Managing Director  
Duke Tech Solutions Inc.  
111 Fieldbrook Ct. Cary, NC 27519  
+1-919-961-6175  
sudheer@duketechsolutions.com

**Ayarah Dharanikota**  
Business Analyst  
Duke Tech Solutions Inc.  
111 Fieldbrook Ct., Cary, NC 27519  
+1-919-376-5657  
Ayarah.dharanikota@duketechsolutions.com

# Table of Contents

| <b>Title</b>                                                 | <b>Page Number</b> |
|--------------------------------------------------------------|--------------------|
| 1. Abstract .....                                            | 3                  |
| 2. Evolution of Telehealth and challenges .....              | 3                  |
| 3. High level stakeholder needs .....                        | 5                  |
| 4. Building blocks of Telehealth Environment .....           | 6                  |
| 5. Cable Operator supported Telehealth recommendations ..... | 9                  |
| 6. Bibliography & References .....                           | 11                 |

## List of Figures

| <b>Title</b>                                                                                                 | <b>Page Number</b> |
|--------------------------------------------------------------------------------------------------------------|--------------------|
| Figure 1 - US healthcare spend 1960-2018.....                                                                | 3                  |
| Figure 2 - Types of Telehealth and relation between Telehealth, Telemedicine and Telecare .....              | 4                  |
| Figure 3 - The Telehealth ecosystem consists of patient, provider and payer side partners .....              | 5                  |
| Figure 4 - High level needs and offer alignment of a Telehealth ecosystem.....                               | 6                  |
| Figure 5 - Day in the life of a patient in different Telehealth scenarios .....                              | 7                  |
| Figure 6 - Telehealth environment incentivized for different ecosystem players .....                         | 8                  |
| Figure 7 - Cable operator high level SWOT analysis .....                                                     | 9                  |
| Figure 8 - A potential roadmap prioritization for executing Telehealth solutions by the cable operators..... | 9                  |


## 1. Abstract

Cable and healthcare industries are crossing paths at many places including in the area of Telehealth. This paper explores how mutually beneficial collaborations can be created between the two industries.

In the process we evaluate different building blocks of Telehealth. A case will be made for the cable operators supporting some of these building blocks. We will provide a quick survey of some of the publicly known Telehealth inter-industry collaborations. We will analyze different stakeholders in the Telehealth ecosystem and their needs.

We are going to highlight the low hanging fruits in this collaboration where Cable operators can begin. And finally provide a preliminary recommendation on how cable operators to venture into a mutually beneficial opportunity.

## 2. Evolution of Telehealth and challenges



**Figure 1 - US healthcare spend 1960-2018**

Healthcare expenses are skyrocketing in United States [1]. Per Capita national expenditure rose from \$146 in 1960 to \$11,160 in 2018 (refer to Figure 1) at a rate of 4.6%. It is expected to rise at a yearly rate of 5.4% from 2019 – 2028 (refer to [2]).

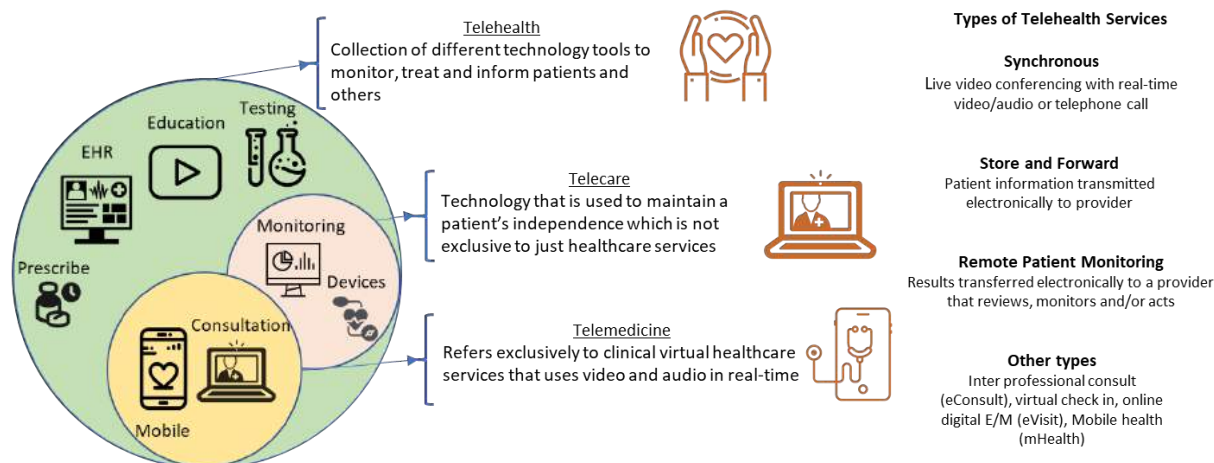
In the market, there are many healthcare and technology initiatives to reduce wastage, reduce cost and increase productivity [3]. As will be discussed later in the document, many inefficiencies can be addressed by Telehealth initiatives. With this brief motivation, we will discuss the background of Telehealth and its promises in detail in the rest of the section.

Although, this is not a primer on healthcare definitions, it is essential to understand tele terms that are often confused by non-healthcare professionals. Telehealth and Telemedicine are defined by HIMSS as [4] – Telehealth: *A broad variety of technologies and tactics to deliver virtual medical, health, and education services. A collection of means to enhance care and education delivery. This term encompasses the concept of “telemedicine,” which refers to traditional clinical diagnosis and monitoring delivered by technology. The term “telehealth”*

*covers a wide range of diagnosis, management, education, and other related healthcare fields including but not limited to dentistry, counseling, physical and occupational therapy, home health, chronic disease monitoring and management, and consumer and professional education.* Figure 2 gives a high-level summary of the above terminology and different types of Telehealth services offered. For a bit more discussion and relevant references from Telecom operators’ point of view refer to [5]. In this paper we focus only on Telehealth aspects of the telecom and healthcare inter industry puzzle.

Telehealth [6], in a nutshell,

- Addresses lack of access to quality care specifically in the remote communities
- Reduces provider burnout in unnecessary activities such as traveling between facilities



**Figure 2 - Types of Telehealth and relation between Telehealth, Telemedicine and Telecare**

- Assists with care gaps in chronic disease management with different patient's providers
- Increases and manages hospital staff utilization through effective capacity management
- Controls the escalating costs the providers are caught in the fee schedules and no-shows

The question is, if Telehealth is solving inefficiencies in the healthcare industry, why is it not taking off? There are some significant technological, policy and perception issues that need to be addressed for Telehealth to be widely adopted. These are briefly mentioned in the insert – “Telehealth Challenges.”

These barriers are, atleast temporarily, lifted due to COVID-19. Radipidly in the last few months Medicare, who is the largest government sponsored health insurance, relaxed their restrictions around the communication mechanisms, payment restrictions etc. [8], [9], [10]. The Telehealth community has gained their long deserved recognition in a hurry. Telehealth solutions or services adoption has surged from roughly 54 percent in 2014 to 85 percent in 2019, indicating a higher level of acceptance and desire for telehealth solutions and services [11]. We will need to wait and see if these policies and payment structures will be continued after the COVID-19 scare subsides. A generic agreement in the healthcare community is that we will not go back to the pre-COVID status, but will also not have the same level of policy and payment relaxation.

Even before the COVID situation itself, there had been significant tailwinds that are driving the innovation and adoption in the Telehealth industry [12]. These tailwinds are mainly around expanded funding availability, increased adoption, proven results, disruptive acquisitions [13], availability of 5G and 10G [14] networks.

### Telehealth Challenges

Telehealth is envisioned in 1960s. Why did it not take off yet, if the healthcare industry unanimously agrees to its benefits? We have three roadblocks **Error! Reference source not found.** -

*Payment:* Payment codes are not available for many Telehealth services yet

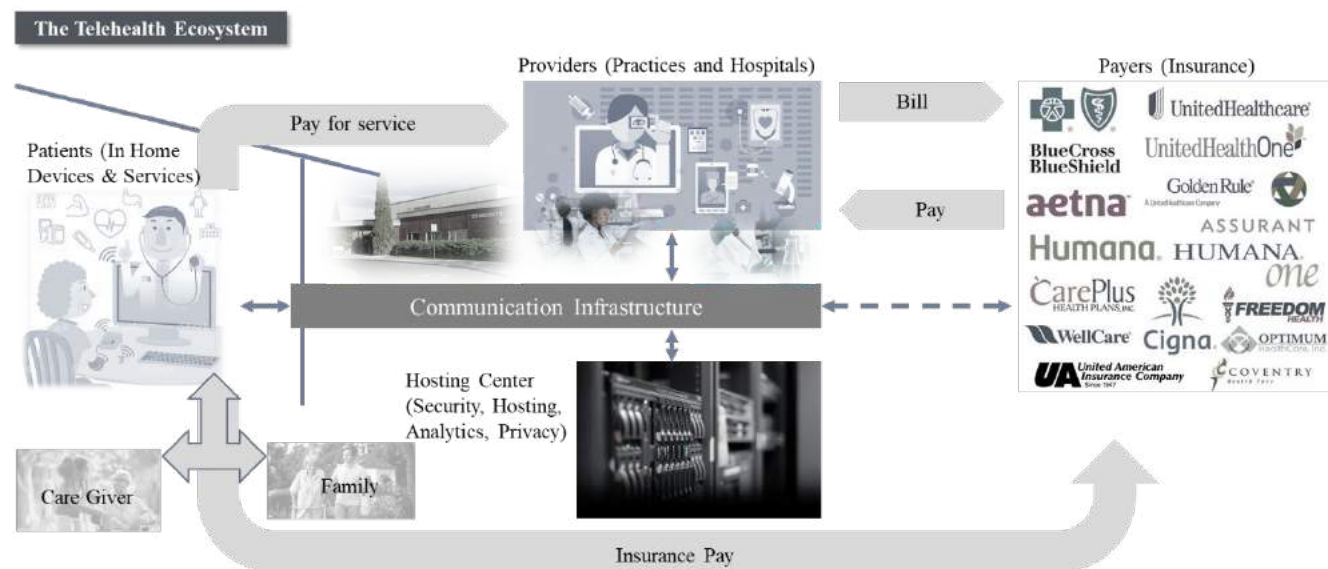
*Policy:* Federal and state healthcare agencies have different policies that restricts the adoption of Telehealth

*Proof of Quality:* Healthcare professionals must prove that the quality of care provided by Telehealth matches non-Telehealth

Telehealth focus has been catching attention from different industries. Firms are trying to solve a portion of the problems. For example, there are over 300 companies claim that they are addressing Telehealth needs [12]. This early adoption stage of the Telehealth industry is heavily fragmented due excitement in the market. In addition, the government agencies, such as The U.S. Department of Health and Human Services (HHS), are trying to control these disruptions in an orderly fashion through relevant Telehealth success metrics [15]. Such a ripe innovation grounds are becoming more and more favorable to the Cable Industry, as we will discuss in the next few sections. The question will be more on what strategy an operator needs to employ when addressing Telehealth needs rather than if they should enter this market.

### 3. High level stakeholder needs

As shown in Figure 3, Telehealth ecosystem has many stakeholders. Of course, it begins with the *patient* and their *family members* that are involved in the decision making of a healthcare visit. If the patient needs additional care, such as taking care of an elderly person, the *care giver* will also need to be part of the ecosystem.



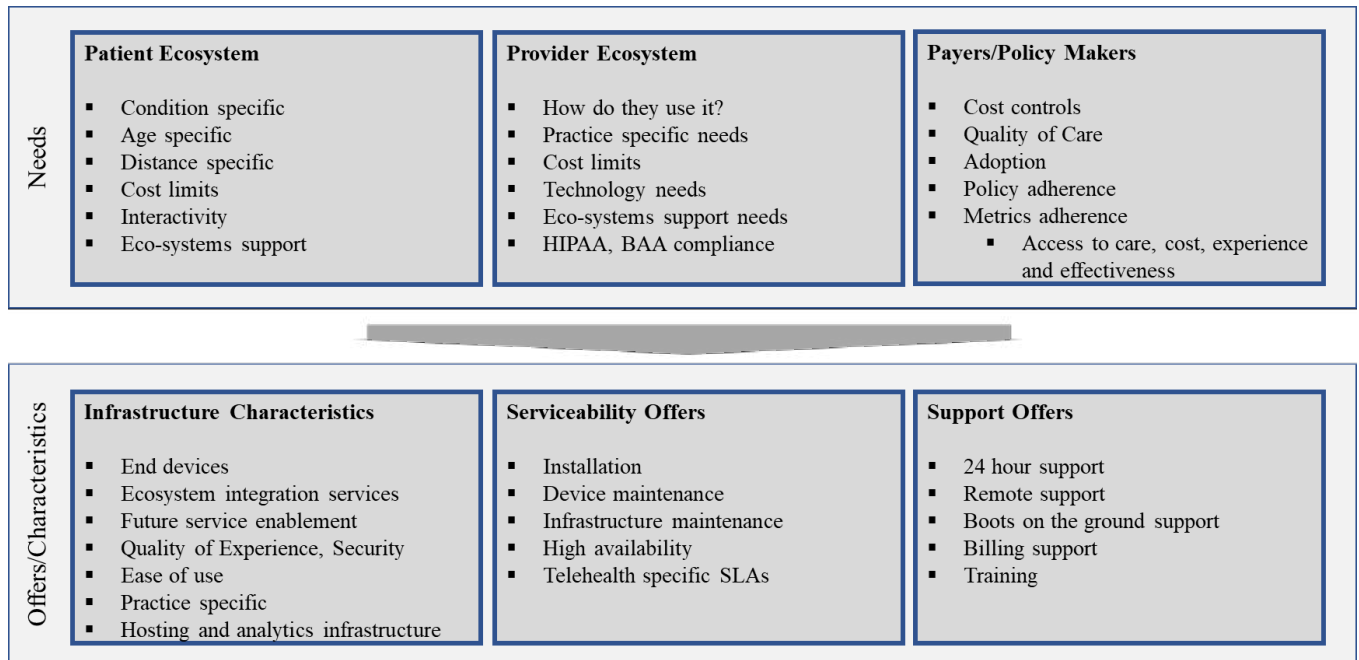
**Figure 3 - The Telehealth ecosystem consists of patient, provider and payer side partners**

The patient interacts with the *provider(s)* (Doctors, Physician Assistants, CNP, etc.) for guidance and treatment. Note, that there could be multiple providers as part of their treatment. In addition, the provider ecosystem includes their *satellite* or *community hospitals* (such as the VA Telehealth services [16]), *pharmacies*, *labs*, and/or *imaging group*. The providers and the patients interact with the *payers* (the public or private insurance companies) for billing and getting paid for the services. In addition, the *state and federal policy makers* (not shown in the figures) sets the guidance to the whole ecosystem. Note, in case of Telehealth, all these interactions are happening on the infrastructure provided by the *operators*. Before diving too deep into the details, let's understand the goals of Telehealth. The policy makers, who are the gate keepers for the Telehealth services, are interested in understanding the efficacy of the Telehealth services in the following dimensions [15] –

- **Access to care:** Access to information for patient, family, care team, and caregiver
- **Financial impact:** Cost to patient, family, caregiver team, society, and the health system
- **Experience:** Patient, family caregiver, care team member, and community experience



- **Effectiveness:** System, clinical, operational, and technical effectiveness



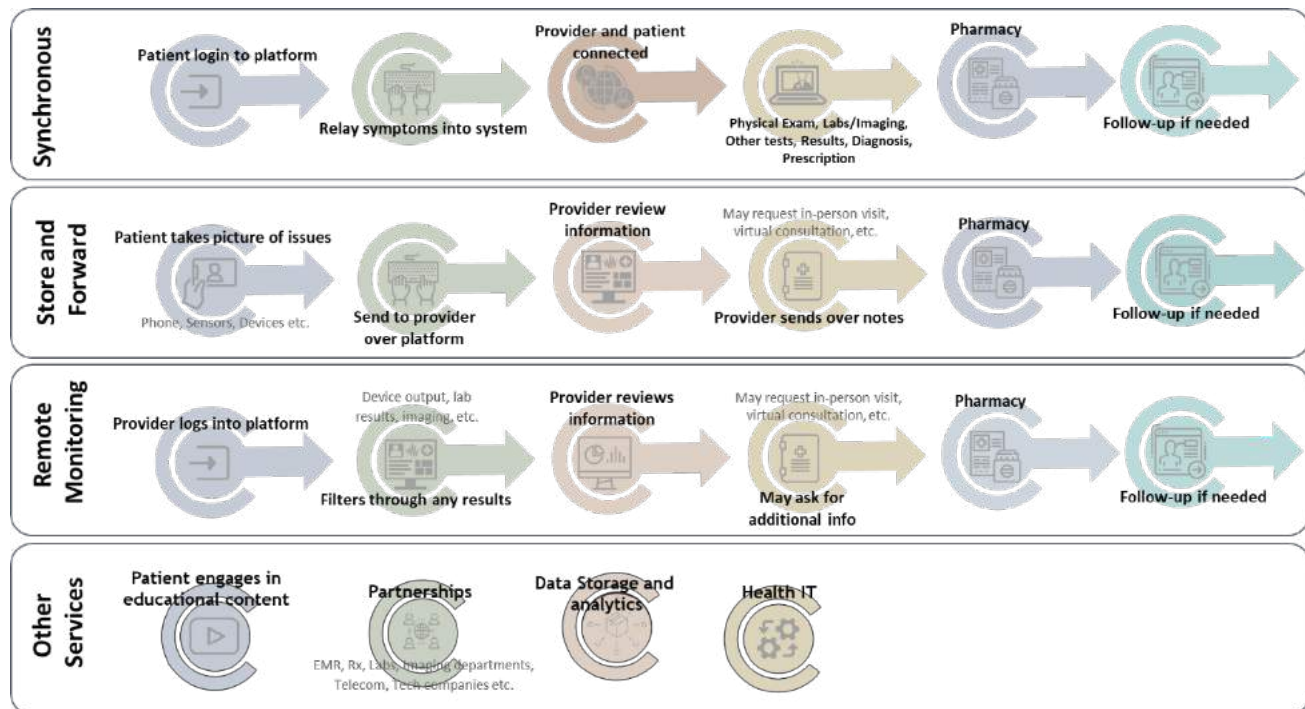
**Figure 4 - High level needs and offer alignment of a Telehealth ecosystem**

Not going too much into the details, we would like to highlight some of the Telehealth needs as shown in Figure 4. These needs can be classified into patient ecosystem needs, provider ecosystem needs, payer needs and policy maker guidance. These needs are self-explanatory and aligns with the Telehealth metrics. To meet the needs of the Telehealth environment, in addition to the most important aspect of the health care involvement, the supporting solution characteristics can be classified into the infrastructure characteristics, the serviceability offers, and the support offers. These are essential to turn a fragmented Telehealth offer into a manageable solution. In the following sections we will build a case on how cable operators can, with their current capabilities, become the Telehealth environment managers.

## 4. Building blocks of Telehealth Environment

The Telehealth environment can be offered as a B2C (Business to Consumer), B2B, and B2B2C business models. As a preliminary step towards offering the Telehealth environment, the Cable operators should investigate which business model they would enter with to ensure success in this highly fragmented healthcare market. Cable operators have a unique relationship with customers through their broadband services. But keep in mind, healthcare is a more emotional and demanding (from service agreements point of view) service. If you are selling to the Telehealth environment through providers and their satellite offices, we need to first consider what are their incentives. These incentives can be identified through the Telehealth metrics identified before. The other more attractive model is to reach consumer through healthcare establishments. There have been multiple studies performed showing this as a valid approach for entering into the Telehealth environment. Many operators have created their healthcare IT initiatives ([17], [18], [19], [20]) to get their feet wet in this inter-industry activity. Lately we have seen the Telco and Cable operators are turning their head to the more profitable Telehealth initiatives (Refer to [22], [23], [24]).

The best way to understand the building blocks of the Telehealth solution is by looking at the day in the life of a patient. As shown in Figure 2, Telehealth offers are classified into Synchronous, Store and Forward, Remote Patient Monitoring and Other categories (Tele-education, metrics gathering, analytics, mHelath, eHealth etc.). These are elaborated in Figure 5 from a patient's point of view to gain more insights into the Telehealth service offerings.



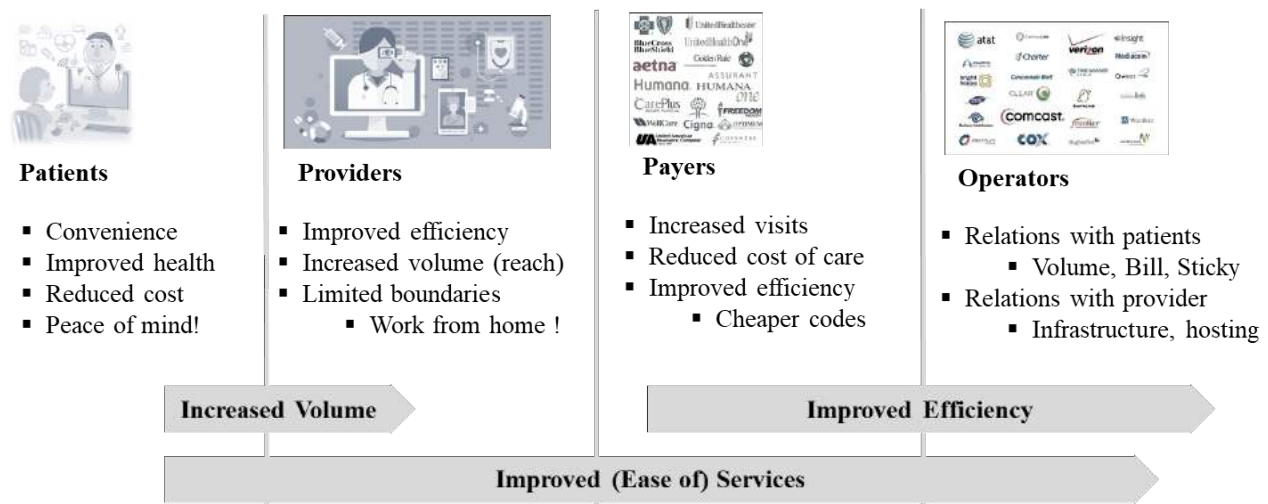
**Figure 5 - Day in the life of a patient in different Telehealth scenarios**

Telehealth is a rapidly evolving industry. The question is how can Cable operators get onto this fast moving train? Looking at the Telehealth environment, as shown in Figure 5, from Cable Operators point of view, one can identify three sets of building blocks as below:

- Patient environment support
  - Communication platform
  - (Future) Consumer oriented unified communications
  - (Future) Device Integration of current and future in home care devices
  - (Future) Usability and accessibility
  - (Future) Supporting team integration
  - (Future) Installation services
  - (Future) Support services
  - (Future) Training support
  - (Future) Build homes to meet future needs
- Provider environment support
  - Communication platform
  - Infrastructure support
  - Unified communications

- Secure infrastructure
- (Future) Integrating the provider ecosystem partners
- (Future) Patient monitoring infrastructure
- (Future) Integrate with store and forward components
- (Future) Increased installation and services relationships
- (Future) Training and education support
- (Future) Telehealth related ML algorithms
- Policy and payer environment support
  - (Future) Telehealth metrics tracking
  - (Future) Telehealth ML algorithms for problem centric analysis
  - (Future) Policy organizational support

Once we understand the business model and the building blocks of services to be offered in the Telehealth environment, we need to address the incentives to get the right priorities for the stakeholders. This process will refine the roadmap of Telehealth environment execution by the cable operators.



**Figure 6 - Telehealth environment incentivized for different ecosystem players**

Figure 6 provides a quick capture of different incentives in the Telehealth environment. When the cable operators are creating their solutions, we recommend elaborating their offering in a framework, such as above, to address the incentives of different stakeholders in the environment. In the following section we evaluate such a framework to create high level recommendations and a roadmap for the cable operators.

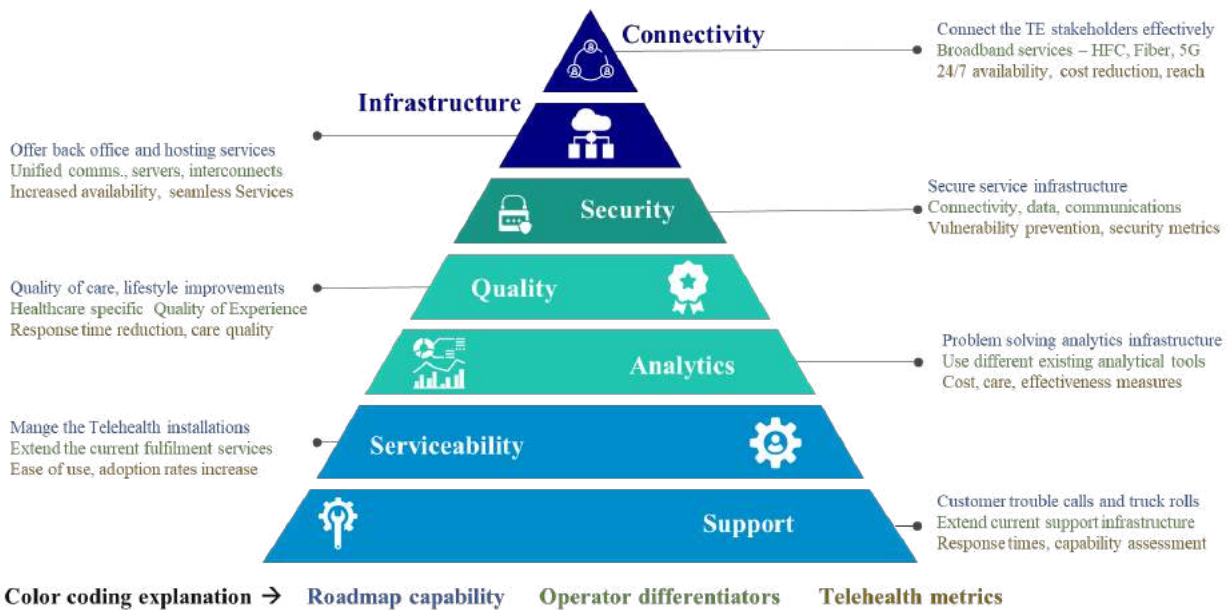
## 5. Cable Operator supported Telehealth recommendations



**Figure 7 - Cable operator high level SWOT analysis**

Figure 7 shows the top differentiators that a cable operator can use to their advantage in this rapidly evolving Telehealth industry, and the risks they need to overcome. To become a formidable player in this inter-industry activity they need to create a razor sharp strategy enhancing their strengths to meet healthcare needs, fostering targeted

relations, and offering aggressive enhanced services with commitment.



**Figure 8 - A potential roadmap prioritization for executing Telehealth solutions by the cable operators**

Bringing all the concepts, metrics, and capabilities together, in Figure 8, we propose a potential roadmap for the Cable Operators to successfully offer a Telehealth environment.

1. Offer your existing capabilities to healthcare industry:
  - a. We propose the operators to play to their strength to start with. This include the *ubiquity of reach and relations* through **connectivity**. It is observed that the biggest challenge for the Telehealth customers is still the broadband connectivity. This is certainly understandable in the rural areas. Even in the urban and suburban areas, this seem to be one of the main concerns. The lowest hanging fruit for the cable operators is to bundle these connectivity services in the healthcare language to address the customer needs. The wireline services such as HFC and Fiber based solutions, along with the wireline services with the future 5G etc., can solve these basic needs. Some of the metrics the healthcare industry understand related to connectivity are the *24/7 availability, cost reduction through efficient use of time and increased reach of the patients*.
  - b. Extend the existing **service oriented infrastructure** to meet the healthcare needs. These include the managed back office and hosting services – specifically focused on the supporting the above mentioned communication relations. This includes services such as Unified communication services. Assist the healthcare providers in hosting their important data (such as EMR data, patient specific data, billing information etc.) and a portfolio of seamless interconnection (between the stakeholders) services. These services can be measured through the *availability and seamless access* metrics.
  - c. Offer a state-of-the-art **secure platform**. The healthcare industry is longing for the day-to-day security infrastructure that the cable operators offer for connectivity, data, and communications. Working with their security infrastructure, the providers by solving the security related issues will be welcome by the healthcare industry. Such an infrastructure’s effectiveness can be measured through *vulnerability prevention and other security metrics*.
2. Adopt your capabilities to the healthcare needs:
  - a. Adopt the service **quality metrics** that the cable operators are using to monitoring to the healthcare services. Develop healthcare specific quality of care and lifestyle improving service – like what we call the Quality of Experience metrics for the services we offer for the triple play services. Metrics such as *response time improvements, cost reduction, quality of care improvements* etc. need to be measured on the data that is mined for these healthcare services. This increases the adoption of the Telehealth services and hence the Cable Operator supported healthcare services.
  - b. Use your **service oriented analytical platform** to assist the complex healthcare issues. Put the complex digital infrastructure that Cable Operators have developed to solve Telehealth related problems. These metrics will be in cost of care reduction, quality of care improvements, Telehealth effectiveness etc. This, in our opinion, is a simple redirection of the analytical infrastructure to the healthcare industry.
3. Increase the capabilities of the operators to meet the future needs of the healthcare industry:
  - a. Develop **Telehealth installation services** as a first step to turn the Telehealth as a standard portfolio service. Extend the fulfilment PPTs (people, processes, and tools) to offer the Telehealth installation services. Implement different fulfillment learning that you have, such as self-service and assisted service combinations, to make the customer’s life easy when deploying these services. Measure your stakeholders and your successes through metrics such as *ease of use and adoption increase*.
  - b. Offer **Telehealth support services** to turn the fragmented market to your advantage. Use your customer’s support infrastructure through care centers, truck rolls to address their healthcare needs. This comes at the expense of mobilizing your support organizations to gain healthcare expertise. The size of the Telehealth opportunity foreshadows the complexities reshaping your service organization. The reward for the operators is significant enough that this is a necessary step to gain the full control of your inter-

industry opportunities. Success can be measured by *response time*, and *problem solving capability assessment* metrics.

In addition to the step-by-step Telehealth services, the Cable Operators have to make the appropriate decisions to develop a go-to-market strategy either through partnerships, building some of the solutions, or by applying the BOT (build, operate and transfer) model. For such a solution they need a clear roadmap for execution, deciding which market they are after: B2B, B2C or B2B2C. For additional information reach out to the authors.

## 6. Bibliography & References

- [1] *The National Health Expenditure Accounts (NHEA) historical data*, available [here](#)
- [2] *National Health Expenditure fact sheet*, available [here](#)
- [3] LeadingAge, *How Telehealth Can Improve Efficiency, Convenience and Outcomes*, June 2014, available [here](#)
- [4] *HIMSS Dictionary of Health Information and Technology Terms, Acronyms and Organizations*, Fifth Edition, available [here](#)
- [5] Ayarah Dharanikota, Sudheer Dharanikota, *Untangling the Tele-X terms for Telecom operators*, August 2020, Duke Tech Solutions blog, available [here](#)
- [6] GlobalMed, *Why Telemedicine, Why Now?* September 2019, available [here](#)
- [7] The role of Telehealth in an evolving health care environment: Workshop summary: Chapter 4 (Challenges in Telehealth), 2012, available [here](#)
- [8] *Medicare Telemedicine health care provider fact sheet*, March 2020, available [here](#)
- [9] *Federal Disaster Resources – Waiver 1135*, available [here](#)
- [10] Susannah Vance Gopalan, *CMS's New COVID-19 Medicare FAQs Provide Detail on FQHCs' Flexibility to Provide Virtual Services During the COVID-19 Emergency*, April 2020, available [here](#)
- [11] mHealth Times, *Definitive Healthcare Survey: Inpatient Telehealth Adoption on the Rise*, August 2019, available [here](#)
- [12] Ziegler white paper, *Deconstructing the Telehealth industry: Part III*, Summer 2020, available [here](#)
- [13] *Best Buy Acquires GreatCall, a Leading Connected Health Services Provider*, August 2018, available [here](#)
- [14] CableLabs, *10G: The Next Great Leap in Broadband*, Summer 2019, available [here](#)
- [15] National Quality Forum, *Creating a Framework to Support Measure Development for Telehealth*, August 2017, available [here](#)
- [16] *VA Tele-Primary Care Hub and Virtual PACT*, 2018, available [here](#)
- [17] *AT&T initiatives on Healthcare IT for digital hospital, caregiver connectivity and connecting pharma etc.*, available [here](#)
- [18] Verizon initiatives on Healthcare IT, available [here](#)
- [19] *U.S. Department of Veterans Affairs Partners with T-Mobile to Help Expand Access to Health Care for Veterans*, available [here](#)
- [20] *CenturyLink Healthcare IT initiatives*, available [here](#)
- [21] *C-Spire Telehealth initiatives*, available [here](#)
- [22] *Spectrum rural Telehealth Solutions*, available [here](#)

- [23] *Comcast joint venture with Quil on digital health initiatives*, available [here](#)
- [24] *Cox Business Emphasizes Commitment to Telehealth through partnership with Trapollo*, available [here](#)

# Cloud-based Dynamic Executable Verification

A Technical Paper prepared for SCTE•ISBE by

**Rafie Shamsaasef**

CommScope  
6450 Sequence Dr, San Diego CA 92121  
858-404-2205  
rafie.shamsaasef@commscope.com

**Aaron Anderson**

CommScope  
117 St. Georges Bay Rd., Parnell  
Auckland, New Zealand  
+64 935 803 75  
aaron.anderson@commscope.com

**Sasha Medvinsky**

CommScope  
6450 Sequence Dr, San Diego CA 92121  
858-404-2367  
sasha.medvinsky@commscope.com



# Table of Contents

| Title                                                                 | Page Number |
|-----------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                  | 4           |
| 2. Secure Boot Example and Limitations .....                          | 4           |
| 3. Software Based Secure Boot Example .....                           | 6           |
| 4. Static vs dynamic analyses .....                                   | 7           |
| 4.1. Static and dynamic analyses concept.....                         | 7           |
| 4.2. Static and dynamic analysis tools .....                          | 8           |
| 5. Tampering Attacks and Threats .....                                | 9           |
| 6. Dynamic executable verification design and concept .....           | 10          |
| 6.1. Related work .....                                               | 10          |
| 6.1.1. Static code signing .....                                      | 10          |
| 6.1.2. Self-checking.....                                             | 11          |
| 6.1.3. Just-in-time code decryption .....                             | 11          |
| 6.1.4. Oblivious hashing.....                                         | 11          |
| 6.1.5. Post-link executable modification.....                         | 12          |
| 6.1.6. Other (intractable) approaches .....                           | 12          |
| 6.2. Goals for integrity protection .....                             | 12          |
| 6.3. Dynamic Executable Verification.....                             | 13          |
| 6.4. Construction .....                                               | 14          |
| 6.4.1. Random function prefixes .....                                 | 14          |
| 6.4.2. Randomly generated check functions.....                        | 16          |
| 6.4.3. Opaque jumtable.....                                           | 16          |
| 6.4.4. Bootstrap.....                                                 | 16          |
| 6.5. Runtime verification.....                                        | 17          |
| 6.6. Security .....                                                   | 17          |
| 6.6.1. Mode 1 .....                                                   | 17          |
| 6.6.2. Mode 2 .....                                                   | 18          |
| 7. Cloud-based architecture for dynamic executable verification ..... | 18          |
| 8. Application use cases .....                                        | 21          |
| 8.1. Browser-based application.....                                   | 21          |
| 8.2. Container-based server application .....                         | 21          |
| 8.3. DRM application .....                                            | 21          |
| 9. Conclusion.....                                                    | 21          |
| Abbreviations .....                                                   | 22          |
| Bibliography and References .....                                     | 22          |

## List of Figures

| Title                                                                                                                                                                                         | Page Number |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| Figure 1 - Example of Linux Secure Boot .....                                                                                                                                                 | 5           |
| Figure 2 – Software Based Secure Boot Example .....                                                                                                                                           | 7           |
| Figure 3 Dynamic executable verification generic use-case.....                                                                                                                                | 13          |
| Figure 4 The DEV module injects random function prefixes (middle), check functions (left), an opaque jumtable (right), and a bootstrap (bottom) into the protected binary at build-time. .... | 14          |
| Figure 5 Dynamic executable verification happens during runtime execution, where each checker function is called according to the mapping defined in the opaque jumtable.....                 | 17          |
| Figure 6 Cloud-based dynamic executable verification .....                                                                                                                                    | 19          |

|                                                                         |    |
|-------------------------------------------------------------------------|----|
| Figure 7 Cloud-based dynamic executable verification (no cloud VM)..... | 20 |
|-------------------------------------------------------------------------|----|

## List of Tables

| <b>Title</b>                                                                      | <b>Page Number</b> |
|-----------------------------------------------------------------------------------|--------------------|
| Table 1 Static and dynamic analysis tools .....                                   | 9                  |
| Table 2 A range of security levels are attainable under different use-cases. .... | 18                 |

## 1. Introduction

Modern software applications are composed of several inner connected modules enabling various features. Today's complex business and market-driven environment constantly pushes the edge to deliver software application faster than ever. Developers are battling with delivery deadlines that are not driven by the complexity of software offerings rather by the go-to-market motivations. As a result, insecure code has become a leading security risk and, increasingly, the leading business risk as well. It's irresponsible at every level to ignore this risk while doubling-down on anti-virus solutions and firewalls — neither of which protects applications [1].

It is important to have holistic view to software protection that provide check points and resolutions throughout the development cycle. It is also equally critical to empower the developers with technologies and methods to be able to automatically identify and detect certain types of attacks. There are commercial software security tools that transform cryptographic credentials so that they cannot be easily extracted. Other tools can make software reverse engineering very hard by sensing a debugger and transforming the binary code logic such that it looks unintelligible even with a debugger attached.

Dynamic Executable Verification (DEV) as described in this paper, provides low-impact dynamic integrity protection to applications that is compatible with standard code signing and verification methods. Further we discuss a system architecture where components of the Dynamic Executable Verification are placed into a secure cloud-based service which can only be configured by an authorized security administrator. To set the context, we discuss secure boot, tampering attacks and methods to perform static and dynamic analyses. Then we dive into details of DEV techniques that aim to ensure that software cannot be tampered with either statically or dynamically, without detection. The cloud aspect of the DEV makes it even easier for developers as the burden of configuring security tools is moved into a cloud service and the risk of releasing an application with lower than intended security is reduced. We will then present a couple of application use cases before concluding the paper.

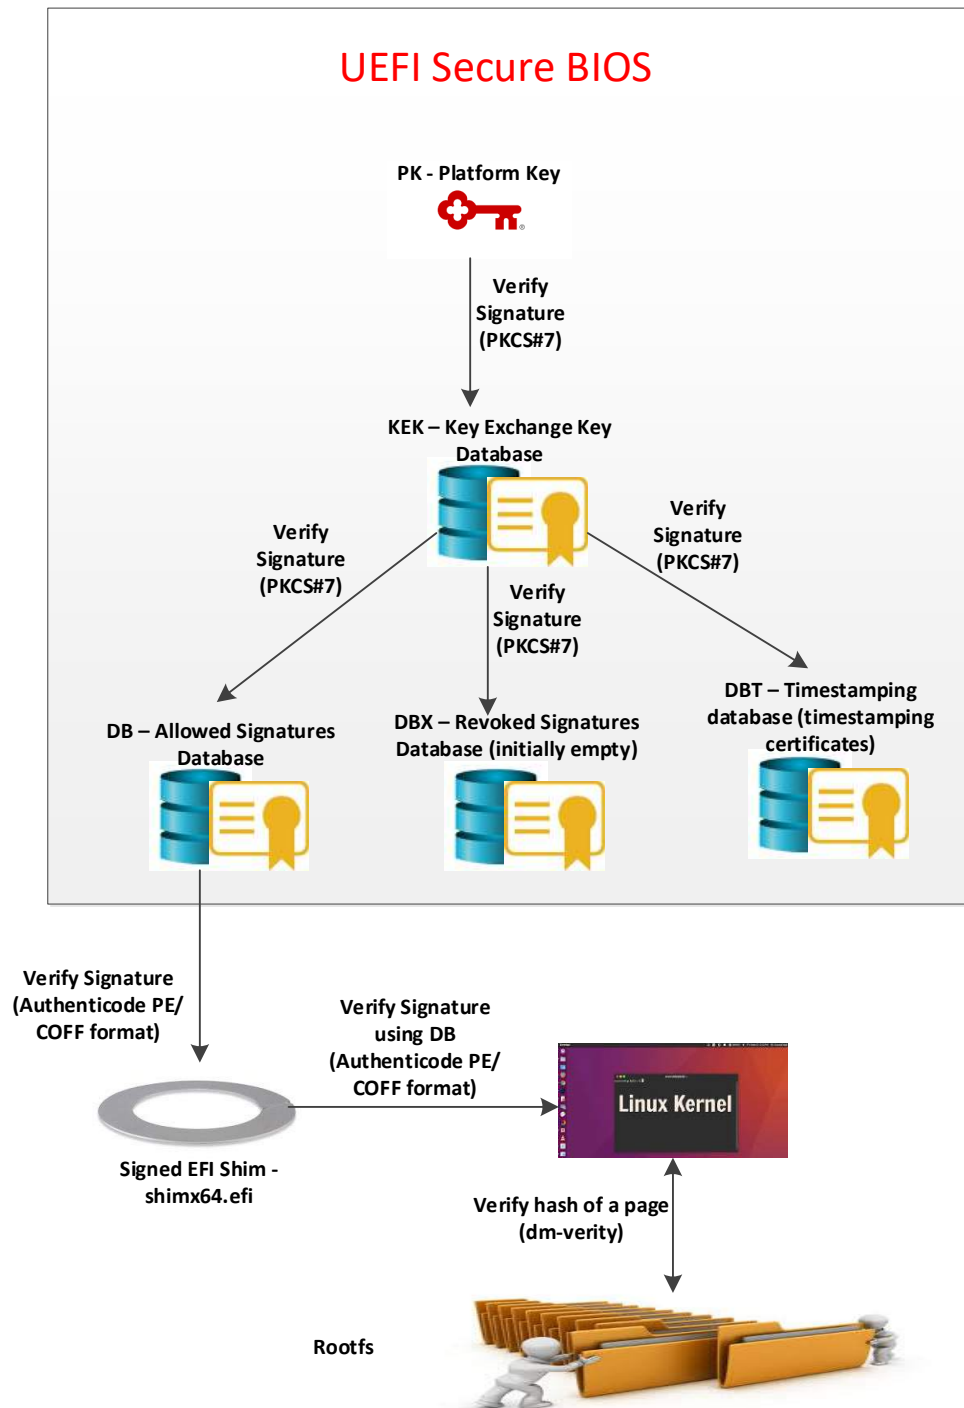
## 2. Secure Boot Example and Limitations

Secure boot implies that each successive stage of software is authenticated – including BIOS or first-stage boot code, successive boot stage, Operating System kernel and all of the applications that execute on top of the OS. Authentication of each of the software layers is repeated every time that a system is re-initialized or rebooted. The very first set of instructions executed by a device after a reboot (first stage boot) may be protected in hardware as read-only (e.g., in ROM or a locked sector or flash) and may not need to be authenticated (since it cannot be changed).

Secure boot prevents physical tampering attacks that include opening up a device and re-flashing all of the software with unauthorized code that has not been digitally signed by an authorized party. If any piece of software is missing a digital signature or if its digital signature is invalid, the device will not boot – or at least that piece of software will fail to execute.

That's quite distinct from secure software download where a software image is authenticated once prior to being persistently installed in the device, into flash or a hard drive. Following a successful secure software download, that software is no longer validated even after a device reboots. A fully secured device would typically include both: secure software download for secure software updates and secure boot to prevent physical tampering with the device.

There are many different ways to achieve secure boot, but one example is illustrated in the figure below:



**Figure 1 - Example of Linux Secure Boot**

In this example, a standard UEFI BIOS that is now commonly available in Windows and Linux PCs as well as embedded devices has UEFI security turned on. It is assumed that UEFI BIOS is HW-protected and cannot be easily modified.<sup>1</sup> The figure shows the standard UEFI key/certificate hierarchy where the top-level Platform Key is used to verify the KEK (Key Exchange Key) which in turn is used to authenticate:

- DB (Allowed Signatures Database) that may for example contain a list of certificates that are permitted to validate a PKCS#7 signature.
- DBX (Revoked Signatures Database) may contain a list of code signatures that are on a prohibited list (e.g., due to security vulnerabilities in the corresponding code)
- DBT (Timestamping Database) – a list of certificates that may be used to validate signed timestamps, utilized to establish when a particular code release was signed. Timestamps may be utilized to reject an older version of the code with an earlier timestamp.

In this example, EFI Shim is first validated by a certificate inside DB and then launched and executed by the BIOS. EFI Shim in turn validates a signature on the Linux kernel (also using a certificate inside the DB) which is subsequently allowed to execute.

This is what is sometimes referred to as UEFI secure boot, but it is incomplete since none of the application code is authenticated and may be easily replaced by an adversary. One way to complete a Linux secure boot (as illustrated above) would be to:

- Place all applications, scripts and executables into a read-only file system such as a Linux rootfs.
- Configure the Linux kernel such that it contains a table of hashes of each page of the rootfs file system, using a facility called dm-verity [2]. Every time that a rootfs page is brought into RAM, its hash is validated. Any tampering to any of the code stored in rootfs will be eventually detected when the corresponding page is brought into RAM and it no longer matches the original hash of that page.

Such secure boot design can be achieved on some devices with custom firmware such as digital set-tops but in some cases, it is not possible to isolate all the executable code into a single read-only file system. Some devices by design have both executable code and data files stored within the same file system, some pages need to be dynamically changed and therefore page-level authentication techniques such as dm-verity do not apply.

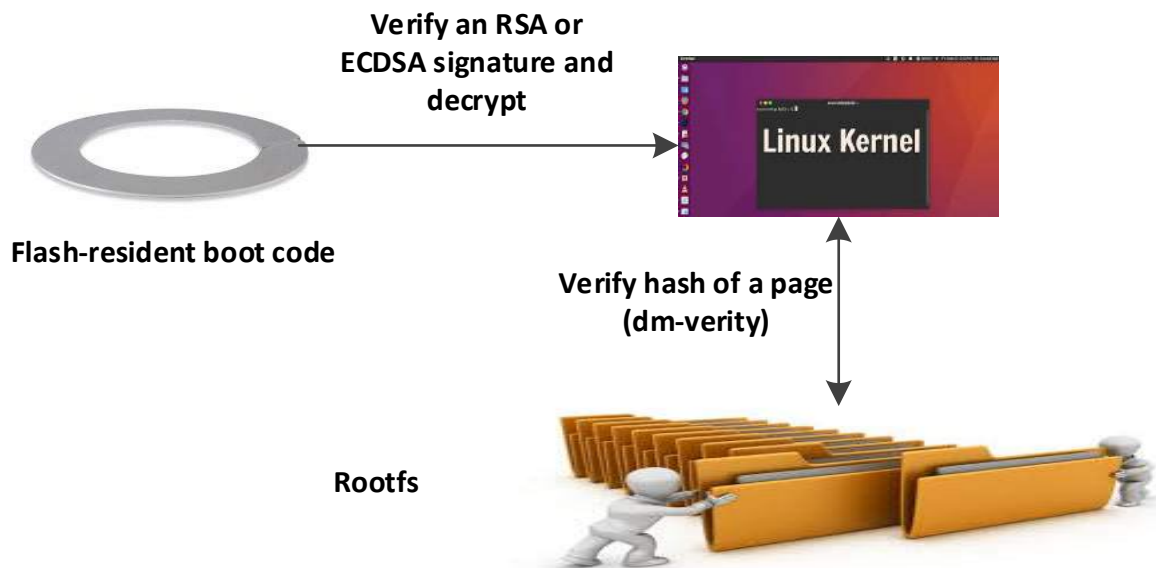
In those cases, one cannot rely on secure boot to validate all of the binary code, especially application code, that may need to execute on a particular device. Techniques described in sections 6 and 7 of this paper provide alternative means to validate the more security-sensitive application code – even in the environments where full or even partial secure boot is not available.

### 3. Software Based Secure Boot Example

This next example of a software-based secure boot is simplified and has less software boot stage than you would normally see, but it illustrates how a software-based secure boot may be applied:

---

<sup>1</sup> Security guidelines for secure UEFI BIOS are provided in NIST SP 800-147



**Figure 2 – Software Based Secure Boot Example**

This device is relying on authenticity of the flash-resident boot code to verify integrity of the Linux kernel which in turn utilizes dm-verity to authenticate all the application code which is isolated in a read-only rootfs file system. This software boot architecture can provide secure boot only if we can be assured that the flash-resident boot code cannot be modified (e.g., by skipping a signature check on the Linux kernel).

Techniques presented in sections 6 and 7 of this paper may be utilized to prevent unauthorized modifications to the flash-resident boot code. Furthermore, this flash-resident boot code may be hiding a decryption key for the Linux Kernel using white box techniques to transform a cryptographic key. An attacker would have a hard time making changes to a self-validating version of the flash-resident boot code. However, replacing it completely would mean that the Linux kernel cannot be decrypted and extracting the decryption key for the Linux kernel is also made hard by white box techniques [3].

## **4. Static vs dynamic analyses**

### **4.1. Static and dynamic analyses concept**

Static and dynamic analyses of a software program are essential ways to discover software security vulnerabilities. They are both important elements of software security analyses and if done correctly, could expose cases that are more costly to resolve once the software is released to market. According to Wikipedia, “Static program analysis is the analysis of computer software that is performed without actually executing programs, in contrast with dynamic analysis, which is analysis performed on programs while they are executing.” [4]

Static analysis requires a good knowledge of software architecture and its internal structures of the software application. It is certainly the more thorough approach and may also prove more cost-efficient with the ability to detect bugs at an early phase of the software development life cycle. It goes deep into the software providing peace of mind that each and every line of source code has been thoroughly

inspected. It potentially can examine all possible execution path without running the program [5]. In a more sophisticated cases, static analysis can be performed on the binary version of the software revealing valuable information about all branches of the code.

Dynamic analysis, on the other hand, is capable of exposing a subtle flaw or vulnerability too complicated for static analysis alone to reveal [6]. It can generate real-time results of execution paths and provides valuable information to security analysts about the software behavior in the runtime environment. In the dynamic analysis, a set of canned input data are typically being supplied to analyze the behavior and expected outputs/results. If done in the controlled and configured settings, it can expose valuable security vulnerabilities in the runtime.

Static and dynamic analyses are only effective if they become part of the software development cycle and performed on every release. This requires certain security policies and practices to be in place per each software release iterations. Policies and practices are evaluated during the design phase from a security point of view. Then static analysis is performed by each developer per each development cycle to identify any leaking security holes. The process continues into the application execution phase where dynamic analysis is conducted in the runtime [3].

As part of this iterative process, developers and security engineers review the analysis report and come up with appropriate recommendations for the development team to consider. The process has to account for a systematic way to feed back the analyses results to the development team and make it a routine practice to adopt. The recommended changes will then be funneled thru the cycle as any other software features or bugs. This systematic and comprehensive process is essential part of any successful static and dynamic analyses practice.

## **4.2. Static and dynamic analysis tools**

The quality and coverage of the analysis dependent on the sophistication of the analysis tools. It varies from those that only consider the behavior of individual source code statements and declarations, to those that include the complete source code of a program in their analysis. Tool can be deployed at unit level, system level or integration level. For a complete list of tools targeting various types of software applications refer to [7] and [8].

A comprehensive evaluation of analysis tool is beyond the scope of this paper. However, it is worth to mention the following criteria in choosing analysis tools:

**Table 1 Static and dynamic analysis tools**

| Criteria                                        | Description                                                       |
|-------------------------------------------------|-------------------------------------------------------------------|
| Programming language                            | C/C++; Java; Perl; Python, etc.                                   |
| Complexity of the code and inner logics         | Code implementing mathematic, crypto and other complex algorithms |
| Input/output data media                         | Ways input and output data flow in and out of the application     |
| Transport, network and communication mechanisms | TCP/IP, HTTP, etc.                                                |
| Targeted Operating Systems                      | Linux, eLinux, Windows, Mac, etc.                                 |
| Build environment                               | Host machine and cross compilers                                  |
| Deployment platforms                            | Devices, servers, browsers, etc.                                  |
| Virtual machines and containers                 | JVM, Docker, etc.                                                 |
| Obfuscation tools                               | Source code level or binary obfuscation                           |
| Configuration scripts                           | Settings and deployment scripts                                   |
| Error handling and logging mechanisms           | How errors and logs are exposed                                   |
| Tool reporting capabilities                     | Reports and results of the analysis tool                          |

## 5. Tampering Attacks and Threats

A consumer electronic device faces two types of general threats:

- 1) Outsider attacks performed by an attacker on the Internet or a public WiFi network. The attacker is for example able to scan the network for vulnerable devices containing outdated versions of the OS, web server, web browser or other commonly utilized applications which have not been recently patched.

Once such a vulnerable device is found, an attacker is able to place her own attack software on that device. Attack software may be utilized to:

- Compromise user's privacy – search user's system for passwords, credit card numbers and other private user information. Even if your own PC and electronic devices are well-protected, your private information may be stolen from one of many online banks, storefronts or other websites where your personal information is exposed. There are numerous and frequent examples of large-scale breaches, including Facebook [9], Yahoo [10], Google [11] and many-many more. And you are not always in control of protecting your personal electronic devices as security flaws may be exploited for a considerable period of time until the security flaw is discovered and patched by the manufacturer. [12]
- Steal user's computing resources for attacker's own use such as Bitcoin mining (called Cryptojacking). User's device will appear to be very slow while much of its memory and CPU resources are utilized by this attacker [13].
- Launch a distributed denial of service attack (after compromising many vulnerable devices) onto popular Internet services or storefronts – for political motives, a personal vendetta or just to create chaos on the Internet. Examples of such attacks that were carried out against GitHub, independent media sites in Hong Kong, CloudFlare security provider and content delivery network, Spamhaus anti-spam service and U.S. banks [14].

Frequent patching of the OS and applications in the devices that you own and setting up perimeter security with firewalls and IP address filtering (including restricting ports and services) would provide a significant deterrent at least against direct attacks against you or your enterprise.



- 2) Insider attacks performed by a subscriber against electronic devices that are owned or leased by the subscriber – in order to rip off digital content, including making illegal copies of movies, songs and games after stripping off or disabling DRM protection.

Additional insider attacks may be performed by disgruntled or dishonest employees, contractors or vendors. Such attacks are generally mitigated with authorization techniques – each person has restricted access only to the computing resources, applications and digital information that are required for that person to perform his or her job.

The focus of this paper is on advanced techniques for protecting software applications against tampering attacks during execution in runtime. Section 6 describes techniques to make software tamper-resistant independent of underlying platform-related code signing protection and without requiring any HW security capabilities on the device where it executes. Section 7 goes further to protect key security parameters that are utilized in the generation of tamper-resistant code on a cloud server with restricted access. Developers are able to utilize a cloud service for creating self-verifying tamper-resistant code without direct access to security-sensitive parameters.

## 6. Dynamic executable verification design and concept

Any software application is potentially vulnerable to *tampering attacks* aimed at defeating their security measures [15] [16] [17]. Tampering attacks are a particularly low effort way to achieve license circumvention and software piracy. In particular buffer-overflow and hooking attacks are common dynamic tampering attacks that regular code-signing methods cannot detect or prevent. The last line of defense (and in many cases the only line of defense) is for applications to be capable of strongly defending their intellectual property and secrets against any such attacks. The goal of integrity protection is to increase the level of effort and complexity of such attacks.

We begin with a brief survey of integrity protection mechanisms in common use. We then describe a novel construction of dynamic executable verification.

### 6.1. Related work

In this section we briefly survey the methods of integrity protection in common usage:

#### 6.1.1. Static code signing

The majority of integrity protection methods in commercial use are targeted at preventing *static tampering attacks*, which involve unauthorized modifications to a program's binary code prior to execution:

- Apple code signing [18].
- Microsoft code signing [19].

#### **6.1.1.1. Drawbacks of static code signing**

Code signing and verification methods do not detect *dynamic* modifications made to the executable code at runtime, such as with *buffer overrun attacks* [20].

### **6.1.2. Self-checking**

Horne et al. [16] and Chang and Atallah [21] present self-checking techniques, in which a program repeatedly checks itself to verify that it has not been modified. These techniques consist of the dynamic (or runtime computation) of a cryptographic hash or a checksum of the instructions in an identified section of code, which is compared with a precomputed hash or checksum value at various points during program execution. Detected tampering will then trigger a response mechanism; such as a silent-failure-mode.

#### **6.1.2.1. Drawbacks of self-checking**

While such methods reliably detect unanticipated changes in the executable code at runtime, it is relatively easy for an attacker to identify the verification routine due to the atypical nature of the operation; since most applications do not read their own code sections [22].

Once detected, these schemes can usually be defeated with simple conditional logic modifications [21] or via hardware attacks [23]; and more recently by *virtual machine debugging attacks* [24], where the address ranges in a program's code section may be translated to an unchanged static image of the code so that any hash or checksum values are always computed correctly despite modifications to the underlying program code.

### **6.1.3. Just-in-time code decryption**

Aucsmith [15] and Wang et al. [25] utilize the notion of self-modifying, self-decrypting code, where any tampering with the encrypted image will result in the decryption of "garbage" instructions, which leads to a catastrophic runtime failure.

#### **6.1.3.1. Drawbacks of just-in-time code decryption**

Several constructions using variations of this technique have been proposed and implemented [26] [27] [28] with varying results; however the widespread adoption of memory protection standards such as PAE/NX/SSE2 [29], and more recently, Intel's MPX [30] with support in mainstream operating systems and toolchains [31], limit this method to legacy and non-standard implementations. For example, since version 18.x of the Microsoft CL compiler/linker, the specification of a writeable attribute on executable code sections is ignored at both compile and in link time.

### **6.1.4. Oblivious hashing**

Chen et al. [17] proposed a technique called *oblivious hashing*, where the idea is to hash an execution trace of a code section. The main goal is to blend the hashing code seamlessly with the code block, making it locally indistinguishable. An oblivious hash is *active* in the sense that the code to be protected must run (or be simulated) in order for the hash to be produced. An oblivious hash also depends on an exact path through a program, as determined by the program's inputs.

#### **6.1.4.1. Drawbacks of oblivious hashing**

Since an oblivious hash depends on a specific control-flow pathway in the executing program, this technique has limited applicability to specialist algorithms with simple linear control-flows. Additionally, since the computation of the oblivious hash is independent of other instructions being executed, tampering attacks that do not affect the internal control-flow of the program, such as *hooking attacks* [32], will remain undetected by this method.

#### **6.1.5. Post-link executable modification**

Many approaches involve the modification of executable code post-linking, so as to inject code or data elements used for the purposes of runtime verification, such as hashes or checksums.

##### **6.1.5.1. Drawbacks of post-link executable modification**

- Incompatibility with standard code-signing and verification methods.
- Limited toolchain compatibility due to possible conflicts with compile-time or link-time optimizations.
- Conflict with technologies that modify binaries post-linking, such as Apple’s *application thinning* [33].
- Potential dependency on external third-party tools to finalize binary representations.

#### **6.1.6. Other (intractable) approaches**

Some methods, such as self-modifying code [34], appear in commercial usage without tractable security descriptions.

##### **6.1.6.1. Drawbacks of intractable approaches**

If the security properties of these methods cannot be objectively assessed, they are unlikely to pass a strict security audit.

### **6.2. Goals for integrity protection**

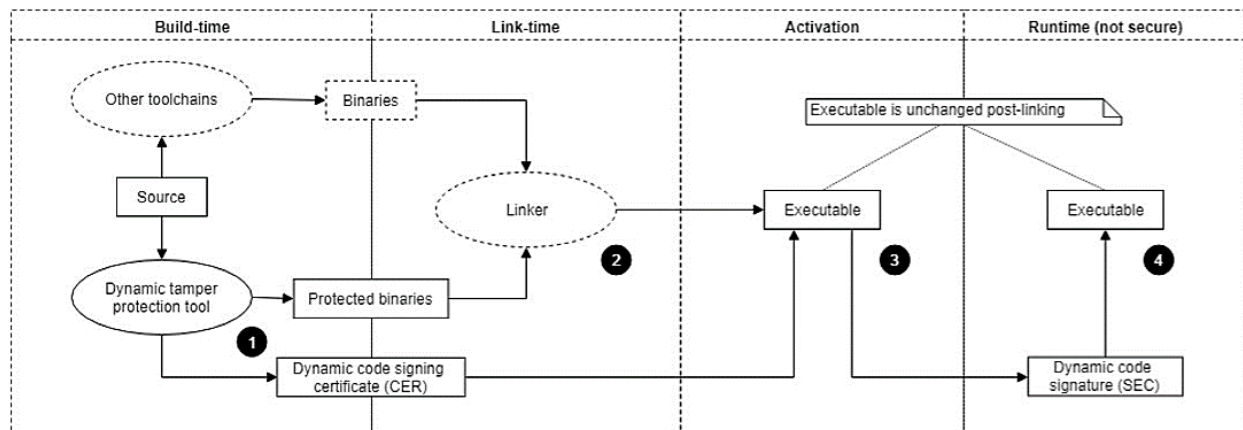
We state the goals of integrity protection that we wish to satisfy with our construction.

1. Increase the cost of static and dynamic tampering attacks
  - Based on tractable principles that can be validated by a security auditor.
  - Resistant to real-world attacks.
  - Based on sound cryptographic principles.
2. Lightweight
  - Low performance overhead.
  - Small memory footprint.
3. Tunable
  - Fine-tunable percentage and locations of automatically generated integrity protection code.
  - Full manual control of integrity protection targets and the locations of verification code.
  - Automatic or custom detection responses.
4. Compatibility with standard code-signing

- Compatibility with standard code-signing and verification methods.
- No frozen/finalized binaries
    - Frozen/finalized binary representations are not required.
  - Support a diverse range of use-cases
    - “Easy mode” where minimizing impact to existing processes is the primary motivation.
    - “Normal mode” for typical use-cases with reasonable expectations of security.
    - “Secure mode” for implementations where security is the primary motivation.
  - Future-proofed formats
    - Use of formats that provide future-proofing against changes to key sizes and cryptographic schemes, such as X.509.
  - Broad platform compatibility
  - Broad toolchain compatibility

### 6.3. Dynamic Executable Verification

In this section we describe a Dynamic Executable Verification (DEV) construction and its security properties.



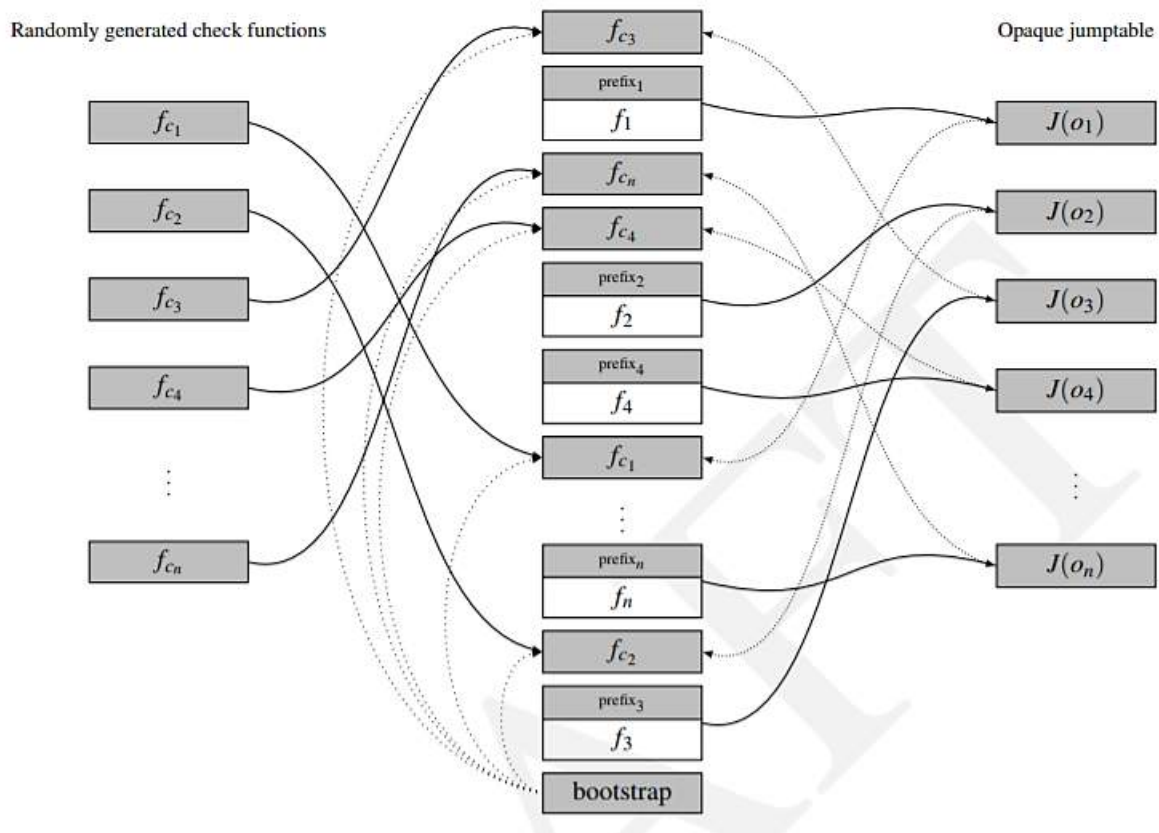
**Figure 3 Dynamic executable verification generic use-case**

- Dynamic executable verification is added to an application at build-time to protect against tampering before and during runtime.
  - The tamper protection tool generates protected binaries along with a dynamic code signing certificate (CER).
  - The tamper detection code in the protected binaries is stealthy, diverse and spread throughout the application to mitigate discovery and circumvention of the tamper protection mechanics.
  - Unlike a static code signing certificate, the CER allows signing of dynamic code blocks within an executable to activate dynamic executable verification of that application.
- Linking is done by a standard linker, which may incorporate unprotected binaries built by other toolchains.
- Activation of dynamic executable verification is achieved by the executable using the CER to self-sign. This should be done in a secure environment to prevent the certificate from being leaked.

- If the certificate is valid and matches the protected portions of the executable, it generates a dynamic code signature (SEC) to be deployed along with the executable to the runtime environment.
  - The executable is unchanged post-linking. This ensures maximum compatibility with sandboxed runtimes and traditional code-signing methods.
4. A runtime failure mode will be triggered if the executable or the SEC has been tampered with in any way before or during execution.

## 6.4. Construction

DEV protection is applied to a program (denoted by the symbol  $P$ ) during compilation by the Dynamic tamper protection toolchain, resulting in a protected program (denoted by the symbol  $P'$ ).



**Figure 4** The DEV module injects random function prefixes (middle), check functions (left), an opaque jumptable (right), and a bootstrap (bottom) into the protected binary at build-time.

### 6.4.1. Random function prefixes

As depicted in Figure 4:

1. For each function/method  $f$  in the input program  $P$ , a random 16 byte function prefix  $f_{\text{prefix}}$  is prepended to  $f$  during compilation, using a standard LLVM operation [35]. We have verified that this method is compatible with all target platforms and toolchains; and with compiler and link-time optimizations.

2. LLVM ensures that the function prefix is aligned to the function's entry point, so that the injected checking code starts from the *prefix* address (typically flagged as an innocuous *data* section) rather than the actual address of the function *f*. This is designed to evade detection by automated and manual analysis techniques to identify self-referential tamper detection code [22].
3. At only 16 bytes per function, the use of a random function prefix is lightweight in terms of footprint, and has negligible impact on performance.

### **6.4.2. Randomly generated check functions**

As depicted in Figure 4:

One or more check functions  $f_c$  (denoting the check function  $c$  of  $f$ ) are randomly generated and injected at random locations in the protected binary  $P'$

### **6.4.3. Opaque jumtable**

As depicted in Figure 4:

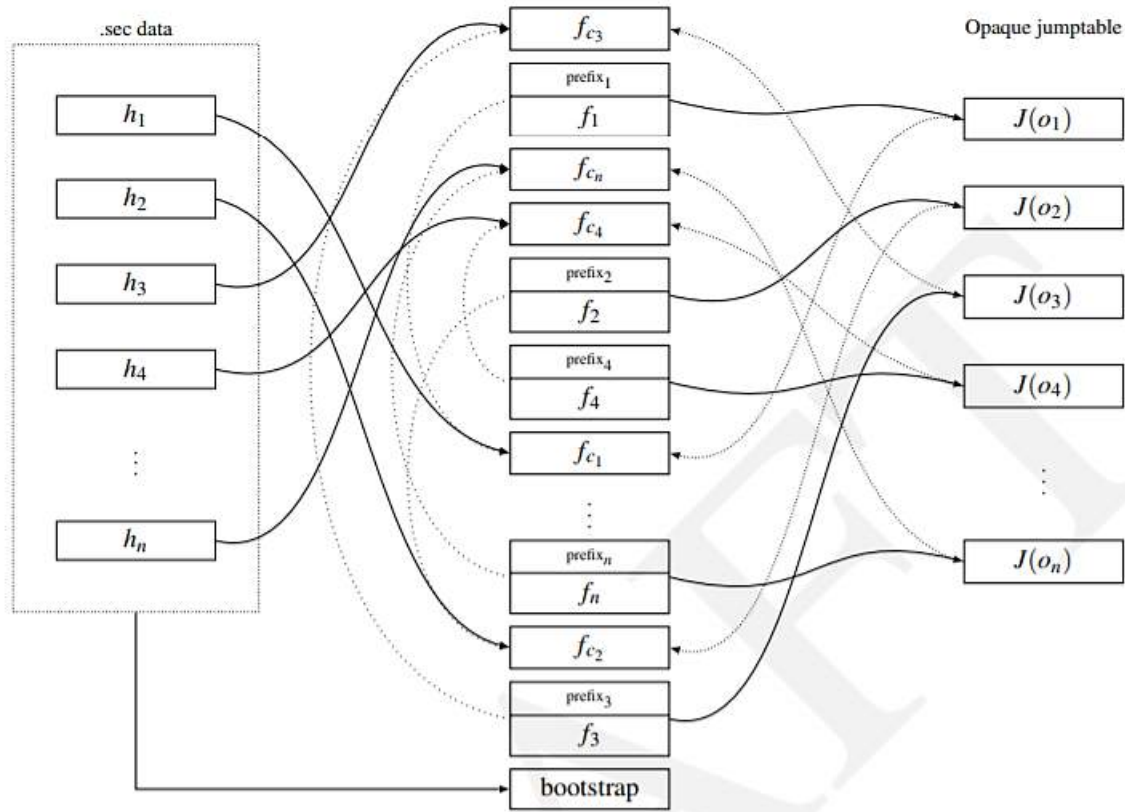
1. The relationship between the calling function  $f$ , and check function  $f_c$  is obfuscated via the use of an opaque jumtable  $J$ .
2. For each randomly generated opaque identifier  $o \in O$ , the opaque jumtable computes the mapping  $J(o) = f_c$  so as to conceal the relationship between the calling function  $f$  and the checking function.
3. Each checking function  $f_c$  references the prefix  $f_{\text{prefix}}$  of  $f$  rather than  $f$  itself, thus avoiding any direct reference to calling function  $f$  from the checking function  $f_c$ .
4. For added security, the mapping  $J(o)$  may be implemented as a complex Boolean expression based on a reduction to a known hard problem, such as 3-CNF-SAT [36].

### **6.4.4. Bootstrap**

As depicted in Figure 4:

1. If the signature (.sec) data does not exist, and if a valid X.509 certificate is passed to the bootstrap, then the bootstrap will *activate* DEV protection by calling each checking function  $f_c \in P'$  (via the opaque jump table  $J$ ) to generate the secure signature data.
2. Activation is carried out one-time per implementation instance.
3. Post activation, the bootstrap reads the .sec data into memory for use during runtime verification.

## 6.5. Runtime verification



**Figure 5 Dynamic executable verification happens during runtime execution, where each checker function is called according to the mapping defined in the opaque jumptable.**

1. After DEV protection has been activated, the .sec data is used at runtime to enforce DEV integrity protection on all specified functions and methods.
2. Detected tampering with the executable or .sec data will result in a failure mode.
  - By default, the failure response should initiate a delayed system crash that is difficult for an attacker to track back to the detection code.
  - This failure mode may be overridden with a custom callback function by specifying the

## 6.6. Security

We have identified two primary security modes that can be utilized in conjunction with the activation method to achieve variable levels of security vs implementation overhead.

### 6.6.1. Mode 1

In mode 1, DEV protection is activated on the *first run* of the protected executable program. DEV protection does not modify this executable at any time. During activation, the DEV bootstrap validates the supplied DEV X.509 certificate (.cer) data. If valid, secure signature (.sec) data is generated, which is used by the executable to enforce integrity protection at runtime.



### 6.6.2. Mode 2

In mode 2, DEV protection is activated by a *setup executable*, which is separate from the runtime executable. DEV protection does not modify either executable at any time. During activation, the DEV bootstrap validates the supplied DEV X.509 certificate (.cer) data. If valid, secure signature (.sec) data is generated, which is used by the runtime executable to enforce integrity protection.

| Security level | Mode | Activation | Runtime   |
|----------------|------|------------|-----------|
| Highest        | 2    | Trusted    | Untrusted |
| High           | 1    | Trusted    | Untrusted |
| Medium         | 1    | Privileged | Untrusted |
| Low            | 1    | Untrusted  | Untrusted |

**Table 2 A range of security levels are attainable under different use-cases.**

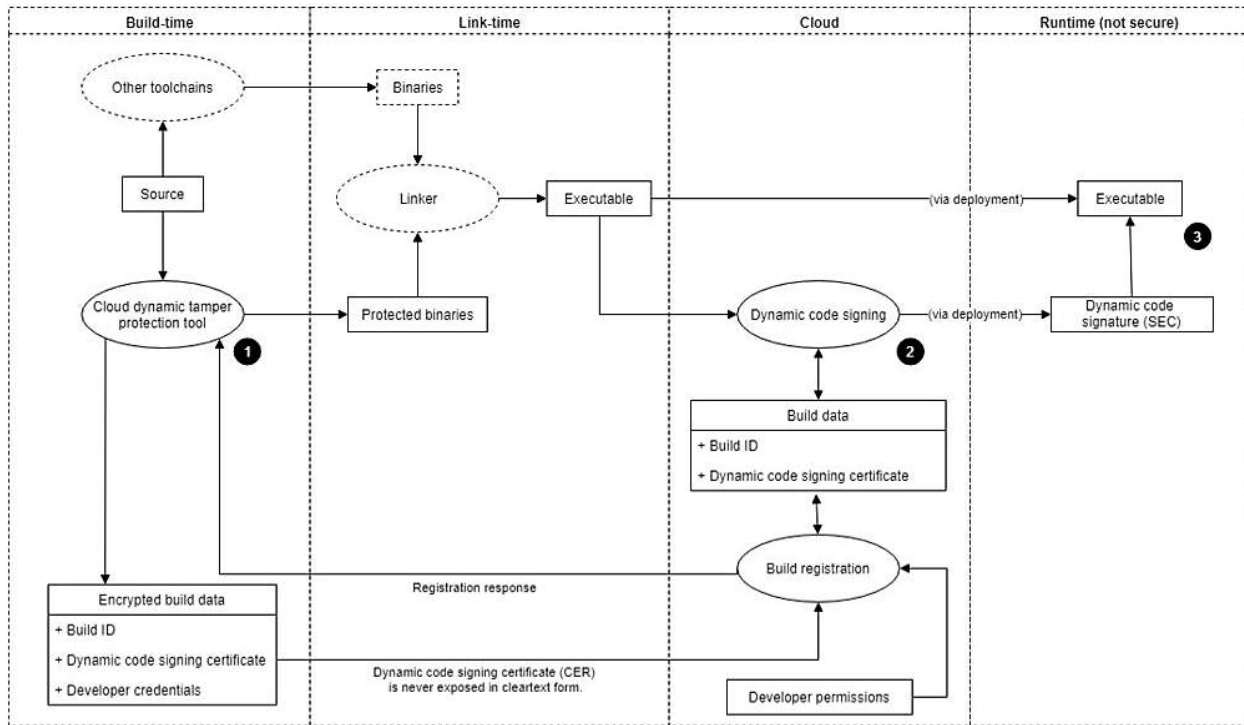
1. Activation in a trusted environment coupled with the splitting of activation vs runtime (in mode 2) provides the highest security level.
2. For high security applications in mode 1, DEV activation should be carried out in a trusted setting, so that the X.509 certificate remains confidential.
3. A medium security level is attainable with mode 1 if DEV activation is carried out in a factory setting or in a privileged runtime mode.
4. A low security profile is obtained in mode 1 if DEV protection is activated *on-the-fly* when the application is first executed at runtime.

## 7. Cloud-based architecture for dynamic executable verification

We present a cloud-based dynamic executable verification architecture to strengthen the overall security properties of DEV. In particular this will offer improved security of the dynamic code signing certificate (CER).

- The cloud service will only permit dynamic code signing request from authorized parties for authorized applications.
- The cloud service consists of two scenarios:
  1. Cloud virtual machine-based activation
  2. Local activation utilizing cloud-based code signing
- This service will be able to interact with other cloud-based security services to offer detailed reporting, metrics, and security alerts.
  - Detailed tamper protection coverage.
  - Activation failures.
  - Runtime metrics.
- Cloud-based dynamic code signing cannot be carried out by developers without cloud credentials or with insufficient permissions.

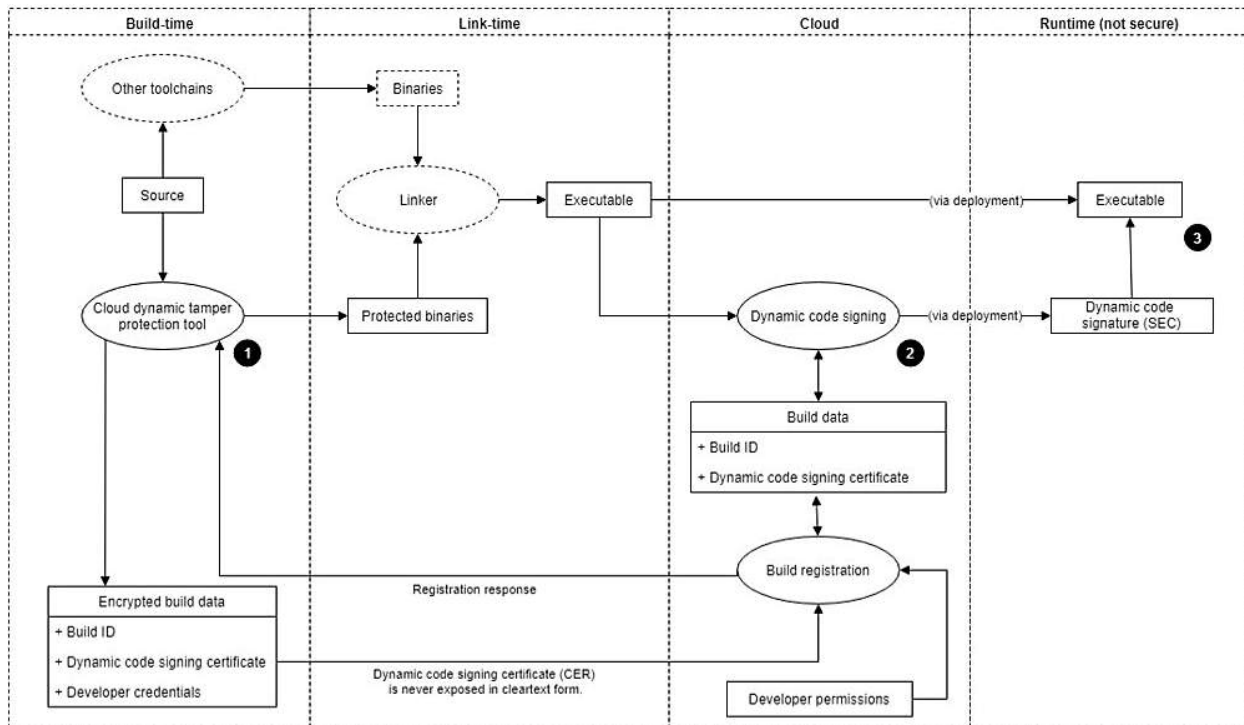
- A correctly defined permissions structure will ensure that only parties with the appropriate credentials can request dynamic signing for production deployment and that signing will only be permitted for applications built with valid developer credentials.
- There is no path to input manually generated CERs from non-cloud tamper protection tools into the cloud-based dynamic code signing. This is by design to prevent unauthorized developers from signing executables for production deployment.



**Figure 6 Cloud-based dynamic executable verification**

1. At build-time, the cloud dynamic executable verification tool sends a unique Build ID and dynamic code signing certificate (CER) securely to the build registration cloud endpoint as depicted in Figure 6.
  - Cloud endpoint authenticates user's credentials before accepting the Build ID and CER over an encrypted link(e.g. using user name/password, digital certificate, one-time password, etc.).
  - Build registration will send a failure response if the request is not authentic or the developer credentials are not authorized for dynamic code signing.
  - Failure responses will abort the tamper protection tool with an error condition.
2. At activation time, the executable is securely sent to a dynamic code signing endpoint post-linking.
  - An authenticated connection to this endpoint is assumed to exist.
  - Cloud service runs executable in a secure VM to obtain a dynamic code signature (SEC).
  - If the executable is invalid, the dynamic code signing fails, or if the developer permissions associated with the Build ID are insufficient, the endpoint will abort with an error condition.
  - Otherwise, the endpoint returns the SEC for runtime deployment.

- Note that the dynamic code signing cloud endpoint does not accept a CER as input, thus preventing non-cloud protected executables from being signed.
3. A failure mode will be triggered if the executable or the SEC has been tampered with in any way before or during execution at runtime.
    - Note that the executable remains unchanged post-linking.



**Figure 7 Cloud-based dynamic executable verification (no cloud VM)**

1. In this scenario, the customer must set up a secure runtime environment to carry out the following.
  - Run a linked executable to generate a Dynamic Code Signing Request (DCSR).
  - Send the DCSR to the dynamic code signing cloud endpoint (via an authenticated connection).
2. The dynamic code signing cloud endpoint evaluates the DCSR.
  - An authenticated connection to this endpoint is assumed to exist.
  - If the request is not authentic, the DCSR is invalid or the developer permissions associated with the Build ID are insufficient, the endpoint will abort with an error condition.
  - Otherwise the endpoint returns a dynamic code signature (SEC) for runtime deployment.
  - As with the base scenario, the executable remains unchanged post-linking.
  - Note that the dynamic code signing cloud endpoint does not accept a CER as input, thus preventing non-cloud protected executables from being signed.

## 8. Application use cases

Dynamic executable verification discussed in this paper opens a great potential with a lot of oversights to designers/developers in protecting their applications. This technology provides ability to delivery software application with self-contained protection against tempering attacks without depending on the targeted platform security offerings. The ability to sign a portion of your software binary and dynamically verifying it in the runtime is applicable to any software application. A wide range of applications can utilize such techniques in their security software development practices and immediately obtain visibility into their application protection. Here we highlight a couple of these use cases.

### 8.1. Browser-based application

Browser-based application is a piece of software running in the browser context. It usually comes in a form of extension (i.e. Chrome) or plug-in (i.e. Edge) or even natively compiled binary (i.e. web assembly in Firefox). As a result, its live time and resource capabilities are limited to the browser session and prone to its attacks. Dynamic executable verification provides independent detection of browser session attacks to the application running in that context no matter which browser is hosting the application.

### 8.2. Container-based server application

As containers become more widespread and acceptable way of deploying server applications, more and more companies are migrating their application to utilize such environment. That of course comes with its own security risk and again dependency on the container provider. As a result, the application security relies on the security of underlaying container technology. Dynamic executable verification provides independent detection of attacks leaking out of hosting containers to applications. This provides a peace of mind to companies virtualizing their application to cloud.

### 8.3. DRM application

DRM agents or libraries hosted in an application can be subject for tampering attacks to lift and later user credentials and authorizations. Dynamic executable verification with cloud-based feature can actively detect and provide visibility to these types of attacks. It gives device/platform independent integrity protection and verification to strengthen the very core features of DRM applications.

## 9. Conclusion

In this paper, we discussed a software protection technology that can fill in the gap in securing applications without any dependencies on the targeted platform. Dynamic executable verification with cloud-based addition helps security engineers and developers to detect tempering threats throughout the code (as needed) in the runtime and take appropriate measures to remedy against them. This enables companies to deploy iterative security analysis processes with independent verifications in their software development practices, utilizing feedback received from attack surface analysis in runtime. There are all kinds of applications in various domains with different use cases that can benefit from this technology and take advantage of its features immediately. CommScope has deployed Dynamic executable verification technique in several products including DRM agents already and in the process of using cloud-based variation in near future.

# Abbreviations

|           |                                                                            |
|-----------|----------------------------------------------------------------------------|
| 3-CNF-SAT | A Boolean satisfiability problem (SAT) in 3-Conjunctive normal form (CNF). |
| CER       | Code signing certificate.                                                  |
| CSR       | Code signing request.                                                      |
| DCSR      | Dynamic code signing request.                                              |
| DEV       | Dynamic executable verification.                                           |
| LLVM      | Refers to the LLVM compiler infrastructure project.                        |
| SEC       | Code signing signature.                                                    |
| UEFI      | Unified Extensible Firmware Interface.                                     |

## Bibliography and References

- [1] Contrast Security, "A Modern Application Security Playbook," 2020. [Online]. Available: <https://www.contrastsecurity.com/8-essential-steps-for-creating-security-strategy>.
- [2] M. Broz, "DMVerity," 2020. [Online]. Available: <https://gitlab.com/cryptsetup/cryptsetup/-/wikis/DMVerity>.
- [3] A. A. Rafie Shamsaasef, "Dynamically Addressing the Gap of Software Application Protection without Hardware Security," in *SCTE Fall Technical Forum*, 2019.
- [4] Wikipedia, "Static program analysis," [Online]. Available: [https://en.wikipedia.org/wiki/Static\\_program\\_analysis](https://en.wikipedia.org/wiki/Static_program_analysis).
- [5] R. Kumar, "Dynamic code analysis VS Static code analysis," 2015. [Online]. Available: <https://www.devopsschool.com/blog/difference-between-dynamic-code-analysis-and-static-code-analysis-2/>.
- [6] N. DuPaul, "Static Testing vs. Dynamic Testing," [Online]. Available: <https://www.veracode.com/blog/2013/12/static-testing-vs-dynamic-testing>.
- [7] Wikipedia, "List of tools for Static Code Analysis," [Online]. Available: [https://en.wikipedia.org/wiki/List\\_of\\_tools\\_for\\_static\\_code\\_analysis](https://en.wikipedia.org/wiki/List_of_tools_for_static_code_analysis).
- [8] S. Parker, "A list of dynamic analysis tools for software," 2019. [Online]. Available: <https://www.peerlyst.com/posts/resource-a-list-of-dynamic-analysis-tools-for-software-susan-parker>.
- [9] A. Glaser, "Another 540 Million Facebook Users' Data Has Been Exposed," Slate, 2019. [Online]. Available: <https://slate.com/technology/2019/04/facebook-data-breach-540-million-users-privacy.html>.
- [10] S. Larson, "Every single Yahoo account was hacked - 3 billion in all," CNN Business, [Online]. Available: <https://money.cnn.com/2017/10/03/technology/business/yahoo-breach-3-billion-accounts/index.html>.
- [11] L. H. Newman, "A New Google+ Blunder Exposed Data From 52.5 Million Users," Wired, 2018. [Online]. Available: <https://www.wired.com/story/google-plus-bug-52-million-users-data-exposed/>.

- [12] Forbes, "Google Shocks 1 Billion iPhone Users With Malicious Hack Warning," [Online]. Available: <https://www.forbes.com/sites/zakdoffman/2019/08/30/google-shocks-1-billion-iphone-users-with-malicious-hack-warning>.
- [13] H. Tuttle, "Cryptojacking: How Hackers Steal Resources to Mine Digital Gold," Risk Management, 2018. [Online]. Available: <http://www.rmmagazine.com/2018/08/01/cryptojacking/>.
- [14] A10, "5 Most Famous DDoS Attacks," [Online]. Available: <https://www.a10networks.com/blog/5-most-famous-ddos-attacks/>.
- [15] D. Aucsmith, "Tamper Resistant Software: Design and Implementation," in *First International Workshop on Information Hiding*, 1996.
- [16] B. Horne, L. Matheson, C. Sheehan and R. Tarjan, "Dynamic Self-Checking Techniques for Improved Tamper Resistance," *Security and Privacy in Digital Rights Management*, pp. 141-159, 2002.
- [17] Y. Chen, R. Venkatesan, M. Cary, R. Pang, S. Sinha and M. Jakubowski, "Oblivious Hashing : A Stealthy Software Integrity Verification Primitive," *5th International Workshop on Information Hiding*, pp. 400-414, 2002.
- [18] Apple Inc, "About Code Signing," 2012. [Online]. Available: <https://developer.apple.com/library/mac/documentation/Security/Conceptual/CodeSigningGuide/Introduction/Introduction.html>.
- [19] Microsoft, "Introduction to Code Signing (Windows)," 2016. [Online]. Available: <https://msdn.microsoft.com/en-us/library/ms537361.aspx>.
- [20] T. Schwarz and S. COEN, "Buffer Overflow Attack," 2004. [Online]. Available: [http://www.cse.scu.edu/~tschwarz/coen152\\_05/Lectures/BufferOverflow.html](http://www.cse.scu.edu/~tschwarz/coen152_05/Lectures/BufferOverflow.html).
- [21] H. Chang and M. J. Atallah, "Protecting software code by guards," *Security and privacy in digital rights management*, pp. 160-175, 2002.
- [22] M. Plasmans, "White-Box Cryptography for Digital Content Protection," 2005.
- [23] P. C. V. Oorschot, A. Somayaji and G. Wurster, "Hardware-assisted circumvention of self-hashing software tamper resistance," *Distribution*, no. June, pp. 1-13, 2005.
- [24] D. Quist and Valsmith, "Covert Debugging Circumventing Software Armoring Techniques," *Blackhat USA 2007 and Defcon 15*, 2007.
- [25] P. Wang, S.-k. Kang and K. Kim, "Tamper Resistant Software Through Dynamic Integrity Checking," *Proc. Symp. on Cryptography and Information Security (SCIS 05)*, 2005.
- [26] J. Cappaert, N. Kisserli, D. Schellekens, B. Preneel and K. Arenberg, "Self-encrypting code to protect against analysis and tampering," in *1st Benelux Workshop on Information and System Security (WISec 2006)*, 2006.
- [27] W. Thompson, "Cryptomorphic programming: a random program concept," *Florida State University, CS Dept., Advanced Cryptography*, vol. 131, pp. 1-11, 2005.
- [28] W. Thompson, A. Yasinsac and J. T. McDonald, "Cryptoprogramming : A Software Tamper Resistant Mechanism Using Runtime Pathway Mappings".
- [29] Hewlett Packard, "Data Execution Prevention," 2005.
- [30] R. Ramakesavan, D. Zimmerman and P. Singaravelu, "Intel ® Memory Protection Extensions ( Intel ® MPX ) Enabling Guide," no. April, 2015.
- [31] Intel Corporation, "Intel MPX support in the GCC compiler - GCC Wiki," [Online]. Available: [https://gcc.gnu.org/wiki/Intel\\_MPX\\_support\\_in\\_the\\_GCC\\_compiler](https://gcc.gnu.org/wiki/Intel_MPX_support_in_the_GCC_compiler).
- [32] S. Vogl, J. Pfoh, T. Kittel and C. Eckert, "Persistent Data-only Malware: Function Hooks without Code," *Network and Distributed System Security Symposium*, no. February, pp. 23-26, 2014.

- [33] Apple Inc, "App Thinning (iOS, tvOS, watchOS)," 2016. [Online]. Available: <https://developer.apple.com/library/ios/documentation/IDEs/Conceptual/AppDistributionGuide/AppThinning/AppThinning.html>.
- [34] G. Tropeano, "Self Modifying Code," *CodeBreakers Magazine*, vol. 1, no. 2, 2006.
- [35] U. o. I. a. Urbana-Champaign, "LLVM: llvm::Function Class Reference," 2016. [Online]. Available: [http://llvm.org/docs/doxygen/html/classllvm\\_1\\_1Function.html](http://llvm.org/docs/doxygen/html/classllvm_1_1Function.html).
- [36] L. A. Anderson, "A survey of control-flow obfuscation methods used in N-Mesh 2," no. October, 2015.
- [37] H. Xu, Y. Zhou and M. R. Lyu, "N-Version Obfuscation: Impeding Software Tampering Replication with Program Diversity," *arXiv*, vol. arXiv:1506, 2015.
- [38] B. Wyseur, "White-Box Cryptography," no. May, pp. 1-9, 2008.
- [39] G. Wurster, P. C. Van Oorschot and A. Somayaji, "A generic attack on checksumming-based software tamper resistance," *Proceedings - IEEE Symposium on Security and Privacy*, pp. 127-135, 2005.
- [40] H. Wee, "On obfuscating point functions," *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pp. 523-532, 2005.
- [41] C. Wang, J. Hill, J. Knight and J. Davidson, "Software tamper resistance: Obstructing static analysis of programs," *Transform*, pp. 1-18, 2000.
- [42] P. C. van Oorschot, "Revisiting Software Protection," *6th Int. Information Security Conf. (ISC 2003)*, vol. 2851, no. October, pp. 1-13, 2003.
- [43] P. C. Van Oorschot, "Overview – Software Protection," no. October, pp. 0-10, 2003.
- [44] T. Tamboli, "Metamorphic Code Generation from LLVM IR Bytecode," *Thesis*, p. 72, 2013.
- [45] N. Runwal, R. M. Low and M. Stamp, "Opcode graph similarity and metamorphic detection," *Journal in Computer Virology*, vol. 8, no. 1-2, pp. 37-52, 2012.
- [46] T. Ogiso, Y. Sakabe, M. Soshi and A. Miyaji, "Software obfuscation on a theoretical basis and its implementation," *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, Vols. E86-A, no. 1, pp. 176-186, 2003.
- [47] J. Nagra, B. Wyseur and T. Herlea, "Trust Model for Software And Hardware-based TR methods," *RE-TRUST Deliverable D*, vol. 2, 2007.
- [48] G. Myles and C. S. Collberg, "Software watermarking via opaque predicates: Implementation, analysis, and attacks," *Electronic Commerce Research*, vol. 6, no. 2, pp. 155-171, 2006.
- [49] W. Michiels and P. Gorissen, "Mechanism for software tamper resistance: an application of white-box cryptography," in *Proceedings of the 2007 ACM workshop on Digital Rights Management*, 2007.
- [50] T. Kerins and K. Kursawe, "A cautionary note on weak implementations of block ciphers," in *1st Benelux Workshop on Information and System Security (WISec 2006)*, 2006.
- [51] P. Junod, J. Rinaldini, J. Wehrli and J. Michielin, "Obfuscator-LLVM - Software Protection for the Masses," *Proceedings of the {IEEE/ACM} 1st International Workshop on Software Protection, {SPRO'15}, Firenze, Italy, May 19th, 2015*, pp. 3-9, 2015.
- [52] M. Jakubowski, C. Saw and R. Venkatesan, "Tamper-Tolerant Software: Modeling and Implementation," *International Workshop on Security: Advances in Information and Computer Security (IWSEC '09)*, pp. 125-139, 2009.
- [53] S. Drape, *Contents Definitions of Obfuscation Evaluating Obfuscation Tools*, 2009.

- [54] CS 276 Lecture 11\_ One-Way Functions \_ in theory.
- [55] D. Corners, The Rootkit Arsenal, 2009.
- [56] Cloakware Corporation, "Software protection and anti-tamper solutions for hostile environments Cloakware Security Suite Solution Overview".
- [57] T. Brekne, "Encrypted Computation," 2001.
- [58] J. M. Borello and L. Mé, "Code obfuscation techniques for metamorphic viruses," *Journal in Computer Virology*, vol. 4, no. 3, pp. 211-220, 2008.
- [59] D. Baysa, R. M. Low and M. Stamp, "Structural entropy and metamorphic malware," *Journal in Computer Virology*, vol. 9, no. 4, pp. 179-192, 2013.
- [60] L. A. Anderson, "Dynamic Executable Verification (DEV)," 2016.
- [61] University of Illinois at Urbana-Champaign, "The LLVM Compiler Infrastructure. Copyright (c) 2003-2014. All rights reserved".
- [62] University of Illinois at Urbana-Champaign, "Exception Handling in LLVM 3.8," 2016. [Online]. Available: <http://llvm.org/releases/3.8.0/docs/ExceptionHandling.html>.
- [63] University of Illinois at Urbana-Champaign, "Cross-compilation using Clang," 2016. [Online]. Available: <http://llvm.org/releases/3.8.0/tools/clang/docs/CrossCompilation.html>.
- [64] University of Illinois at Urbana-Champaign, "Clang command-line options," 2016. [Online]. Available: <http://llvm.org/releases/3.8.0/tools/clang/docs/UsersManual.html>.
- [65] University of Illinois at Urbana-Champaign, "lrc - LLVM static compiler," 2015. [Online]. Available: <http://llvm.org/docs/CommandGuide/lrc.html>.



# **Addressing Unrelenting Growth In Backbone Fiber Systems Using Next Generation Photonics And Automation**

A Technical Paper prepared for SCTE•ISBE by

**Timothy Maenpaa**  
Consulting Regional Systems Engineer  
Ciena Corporation  
Alpharetta, GA  
(678) 395-3605  
tmaenpaa@ciena.com

# Table of Contents

| <b>Title</b>                                          | <b>Page Number</b> |
|-------------------------------------------------------|--------------------|
| Table of Contents .....                               | 2                  |
| 1. Introduction.....                                  | 3                  |
| 2. Preface .....                                      | 3                  |
| 3. Network Design .....                               | 4                  |
| 3.1. Migration from Older technologies .....          | 4                  |
| 3.2. Express Overlay Networks .....                   | 5                  |
| 4. New Hardware.....                                  | 6                  |
| 4.1. Higher Baud, Higher Bandwidth Transponders ..... | 6                  |
| 4.2. Flexible Grid Photonics .....                    | 7                  |
| 4.3. L-Band Capability .....                          | 8                  |
| 4.4. Additional Fiber Diagnostics.....                | 10                 |
| 5. Software and Automation .....                      | 10                 |
| 5.1. Software – From Planning to Operating.....       | 10                 |
| 5.2. Planning Equals Deployment .....                 | 10                 |
| 5.3. Zero Touch Provisioning .....                    | 11                 |
| 5.4. Margin Mining .....                              | 12                 |
| 6. Conclusion.....                                    | 13                 |
| Abbreviations .....                                   | 14                 |
| Bibliography & References.....                        | 14                 |

## List of Figures

| <b>Title</b>                                                                                                                      | <b>Page Number</b> |
|-----------------------------------------------------------------------------------------------------------------------------------|--------------------|
| Figure 1: 11 Year ‘Full Spectrum System’ Growth .....                                                                             | 4                  |
| Figure 2: Fixed Grid to Flexible Grid Migration .....                                                                             | 5                  |
| Figure 3: Express Overlay Network .....                                                                                           | 6                  |
| Figure 4: Capacity vs. Reach Performance by Baud .....                                                                            | 7                  |
| Figure 5: Flexible Grid Photonic Use Cases .....                                                                                  | 8                  |
| Figure 6: First Generation L-Band Impact from Stimulated Raman Scattering (SRS) and Power<br>Transfer during L-Band Upgrade ..... | 9                  |
| Figure 7: Integrated Amplified Spontaneous Emission (ASE) – Optimized C&L-Band Upgrade .....                                      | 9                  |
| Figure 8: Zero Touch Provisioning .....                                                                                           | 11                 |
| Figure 9: Optimizing capacity for SNR requires programmable optics AND real-time analytics .....                                  | 13                 |

## 1. Introduction


There is a limit to how many bits can be sent down an optical fiber. This limit, known as the ‘Shannon Limit’ is defined as the maximum rate at which data can be sent over a medium with zero errors. Technologies such as coherent optics allow operators to get closer to this theoretical limit. However, moving forward, the gains in spectral efficiency that will be achieved with future generations of coherent technology are diminishing. This will require alternative approaches and ideas to deal with network scalability challenges. Using updated designs, next generation hardware, and software tools, network operators will be able to extend the life of their networks as well as deploy new networks more quickly, efficiently, and accurately. Using these ideas, operators can start reducing the slope of their spectral usage curve while at the same time, deploying new photonic networks with higher efficiency.

## 2. Preface

It is estimated that global internet traffic will grow 3.7-fold from 2017 to 2022. Globally, IP traffic (alone) has also grown three-fold in the same period. This reflects a compound annual growth rate (CAGR) of 30% and 26% respectively<sup>1</sup>. Drivers for this tremendous growth include Video (IP, Internet, VoD), gaming, mobile devices, social media, and the Internet of things. Additionally, the average residential bandwidth speed has more than doubled from 24 Mbps to almost 50 Mbps. Consumers now regularly achieve 100+ Mbps download speeds with “standard” cable internet service.

Looking into the future, will these tremendous growth rates continue? And what are the drivers? The answer to the continued growth question is “yes”. The drivers will be new and expanded offerings that will continue to drive the need for bandwidth such as medical imaging, tele-medicine, virtual reality and gaming, cloud storage, and of course this new necessity of “working from home”. These new drivers, as well as the old ones, will continue to drive demand for the internet and fuel the growth of optical networks.

How much bandwidth will be needed? While this is certainly a “loaded” question, let us consider some simple math that could provide some direction for thought. The C-band, or conventional band, covers the fiber spectrum from 1530 nm to 1565 nm and is approximately 4800 Ghz wide. That might be forty-eight 100 Ghz spaced channels or that might be ninety-six, 50 Ghz spaced channels. If we consider ninety-six 50 Ghz spaced, 35 GBaud, 100 Gbps channels as a “full spectrum system”, how many of these “full spectrum systems” are required if we project a 25% CAGR into the future for the next 11 years?



**Figure 1: 11 Year ‘Full Spectrum System’ Growth**

This graph shows that using a compound annual growth rate of 25% per year, an operator using a photonic layer capable of ninety-six 100 Gbps channels will, by year 7 have to deploy and enable a new, full spectrum, photonic system with EVERY subsequent year. When the provider hits year 11, the growth rate exceeds TWO new, full spectrum systems every year.

So, how are operators to keep up with this growth? This paper has been written to suggest strategies that might be used by multiple-system operators (MSOs) to efficiently scale their networks and support more capacity with fewer incremental systems/fiber pairs (and network resources), and improve efficiencies for their customer as well as for their shareholders.

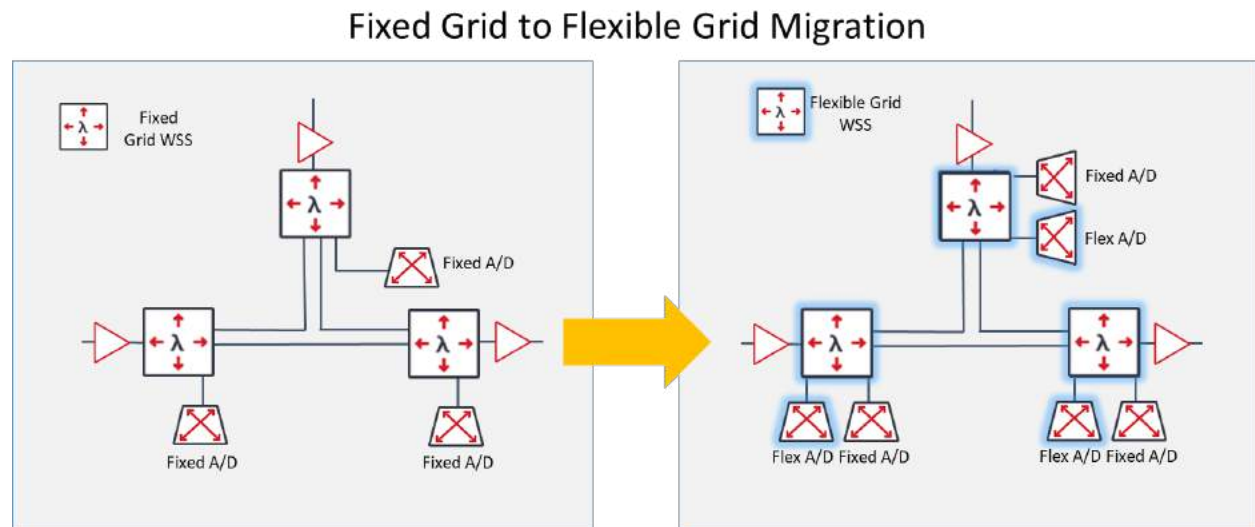
### 3. Network Design

#### 3.1. Migration from Older Technologies

Older generation photonic networks are typically based on 100 Ghz or 50 Ghz ITU grid and use passive, fixed grid filters to provide optical channel access. These older systems are reaching capacity based on the limited spectral efficiency of this generation of photonics and transponders. The first opportunity that operators could use to slow their system deployment curve is to upgrade an older, fixed grid network to a newer, flexible grid photonic system.

Much of today’s newer photonic equipment is still backward compatible with the older, fixed grid standards. This would allow, for example, flexible grid WSS hardware to provide a direct replacement for older fixed grid WSS hardware in the existing fixed grid network. This equipment compatibility would permit network operators to upgrade high use portions of their existing fixed grid system, using the current fiber pair, to newer, flex grid capable wavelength selectable switch (WSS) hardware. Additionally, flex grid capable colorless channel mux/demuxes would also be installed with the new flex grid WSS modules. Once the flex grid capable photonics are in place, flex grid functionality could be “enabled”. Operators could then utilize the full functionality of flex grid. This fixed to flex upgrade does

come with challenges such as migration maintenance windows, specific channel pass-through rules, and operational differences between the fixed and flex grid portions of the network. However, if specific portions of the fixed grid network are bandwidth limited, this type of flexible grid upgrade is an option available that could increase the spectral efficiency by as much as 350% on the new flex grid portion of the network.



**Figure 2: Fixed Grid to Flexible Grid Migration**

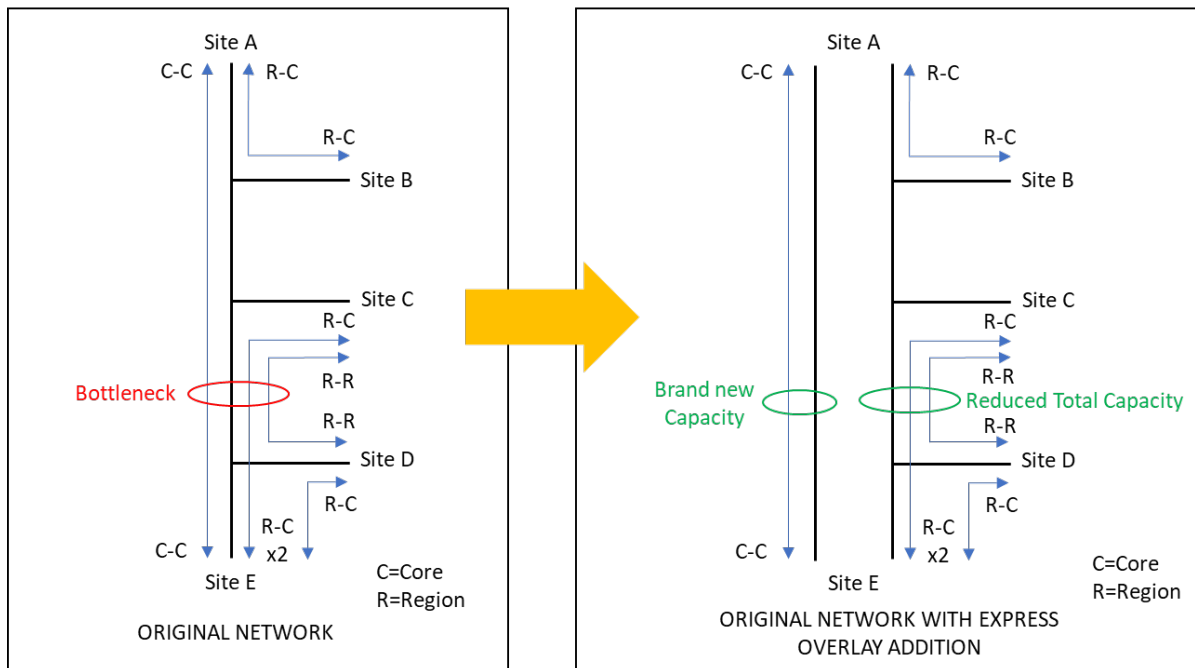
With proper network planning during the fixed to flex migration, the spectrum can also be “de-fragged”. The de-frag would allow for the efficient continued use of the operator’s current transponders while also allocating space for the use of next generation, larger bandwidth transponders. This upgrade enhancement will delay network exhaust because of the higher spectral capacity of the next generation hardware while continuing to utilize existing transponder assets that the operator has already purchased.

### 3.2. Express Overlay Networks

Many current MSO networks were built as “one size fits all” networks. The one size fits all network must provide all transport functions including Core-to-Core, Core-to-Region, and Region-to-Region connectivity. These different functional traffic flows cause the network to grow at different rates. When highly used portions of the network grow faster than other portions, bandwidth bottlenecks can occur.

Single use, express overlay networks is another strategy available to MSOs that could provide targeted relief and extend the life of the original network. In the portions of the network where bandwidth bottlenecks are starting to occur, an express overlay photonic network could be added that would address the highest functional contributor to the bandwidth. As an example, if there is a bandwidth bottleneck in a Core-to-Core corridor, a new, purpose-built Core-to-Core express overlay network could be added. This does two things. First, the original Core-to-Core traffic would be migrated to this new network path. This migration of bandwidth from the original to the new express overlay system would free up spectrum on the original network for additional Region-to-Core and Region-to-Region traffic. This extends the original network’s “time until full”. Second, the new path could be built using new technologies (flexible grid, next generation transponders) that would then increase the spectral efficiency of the new path helping to slow even further the need for additional photonic systems.

## Express Overlay Network



**Figure 3: Express Overlay Network**

## 4. New Hardware


In addition to network design strategies, there is a new generation of photonic equipment that can dramatically increase spectral efficiency and thus increase the time required between network builds. This new hardware includes higher baud, coherent modem technology that creates higher bandwidth channels to optimize reach and eliminate regeneration. Flexible grid photonics facilitate the efficient transport of higher baud transponders and ultimately increase the spectral density of the fiber and lower the cost per bit. L-Band photonics that enable the L-Band spectrum adjacent to the C-band allow for doubling the capacity of the fiber, increasing the return on investment of fiber assets and delaying new fiber builds.

### 4.1. Higher Baud, Higher Bandwidth Transponders

Using today's newest optical transponders, operators can make transponder selections that exactly fit their network transport needs. "I have a short link, and I need maximum fiber capacity." Or, "I have a long link and I don't want any regeneration." And, "I want to optimize my spectral efficiency." Today's newest generation optical transponders can provide a solution for all of these needs.

With the ability to select coherent optic technology with baud ranging from 35 GBaud up to 95 GBaud, spectral efficiency can be optimized for the given photonic layer topology. Today's transponders use probabilistic constellation shaping which creates the ability to select optical transponder line rates from 100 Gbps to 800 Gbps to optimize the capacity relative to available margin. The new transponders also use enhanced forward error correction (FEC) and other sophisticated algorithms that ultimately permit longer reach of the photonic signal.

These next generation transponder features allow network operators the ability to independently select the baud and line rate to maximize the spectral usage across a desired path. As an example, for shorter reach paths an operator would be able to create an 800 Gbps link while using a symbol rate of approximately 90 GBaud. Then, provision a second transponder pair to create a non-regenerated 400 Gbps line using 95 GBaud on a path that goes from Miami to Seattle. Finally, a third transponder pair could be used to “dial in” the optimal bit rate to satisfy any specific route in-between. This degree of transponder flexibility is creating a new paradigm in how photonic networks can optimize their spectral use.



**Figure 4: Capacity vs. Reach Performance by Baud**


## 4.2. Flexible Grid Photonics

Older generation optical networks are based on 100 GHz or 50 GHz spaced photonic systems. These older, gridded networks can offer forty-eight or ninety-six fixed grid optical channels within the total 4800 GHz C-Band spectrum. This fixed grid spacing is based on the ITU standard and has been the norm for most photonic systems for over 20 years. These older systems use passive, ITU grid filters to provide wavelength ingress/egress. Second and third generation optical transponders running at 35 GBaud typically require 37.5 GHz of optical spectrum which fit perfectly into these ITU gridded filters. This version of the network has served the industry well for years.

As next generation transponders become available, they bring with them the need for larger per channel spectrum. This larger channel spectrum requirement exceeds that which is available on these ITU gridded filters. Flexible grid photonics offer bandwidth “chunks” as small as 6.25 GHz. The flexible use of these 6.25 GHz chunks as well as a WSS based mux/demux for wavelength ingress/egress, permits more granular utilization of spectrum needed for the desired network transmission characteristics.

With flexible grid photonics and the ability to change the baud and bit rate of each specific transmitter, network operators can optimize the use of the spectrum. In the situation where the operator provisioned an 800 Gbps channel using a symbol rate of 90 GBaud, this configuration would provide as many as 48 (x800 Gbps) channels which is a total bandwidth of 38.4 Tbps on the fiber pair for a short reach path in the C-Band. If the L-Band is in place, that number is doubled. In the example of a path that crosses the United States running at 400 Gbps and 95 GBaud, this configuration would provide as much as 33.6 Tbps (42x400 Gbps channels) transported coast to coast without any O-E-O regeneration.

Flexible grid photonics add another level of adjustability to the new optical network. With the ability to specify the guard band appropriate for the link design, operators can engineer their photonic deployments with the highest spectral efficiency available for their specific network. For example, on shorter photonic paths, packing multiple network media channels into a single media channel allows for a reduced guard band size which permits the channels within the media channel to occupy less spectrum. This ultimately allows more channels to be placed in the total spectrum. This guard band reduction is another way that flex grid photonics improve spectral efficiency.



**Figure 5: Flexible Grid Photonic Use Cases**

### **4.3. L-Band Capability**


As the incremental gains in spectral efficiency progressively diminish with each new generation of coherent technology, expanding the photonic layer into the L-band is becoming an increasingly popular option for scaling networks. The L-Band, long band, or extended band is the wavelength band immediately next to the C-Band. The L-Band covers the spectrum from 1565 nm to 1625 nm. For years, the L-Band has held the promise of extending the usable spectrum for operators. Unfortunately, first generation L-Band hardware never commercially delivered the additional L-Band capacity due to usability issues and deployment/upgrade complexities.

First generation L-Band deployments required splitter/couplers to provide fiber access to the L-Band equipment that was to be added later. These splitter/couplers used valuable span margin ultimately reducing much needed receiver optical signal to noise ratio (OSNR) on early generation transponders.

Additionally, when the L-Band is added to the C-Band in an active network, the L-Band can experience optical amplification at the expense of the C-Band due to stimulated raman scattering (SRS). So, when the L-Band is finally added to the C-Band fiber, the C-Band could experience a reduction in optical power. This reduction in C-Band total power has a potential impact to the OSNR of those in-service C-Band channels, especially when the C-Band spectrum is full. This link performance challenge is made harder by different fiber types, span losses, raman configurations, and channel powers in the network. As a result, the link engineering for the L-Band addition can be quite complicated.



### First-Generation L-Band Impact from Stimulated Raman Scattering (SRS) and Power Transfer during L-band Upgrade




**Figure 6: First Generation L-Band Impact from Stimulated Raman Scattering (SRS) and Power Transfer during L-Band Upgrade**

Finally, the first-generation L-Band upgrade challenge was further complicated by the fact that none of the L-Band hardware had been added day one. This meant that every site in the network had to be visited to add the needed L-Band equipment. On larger networks, a visit to every site could be very costly and time consuming.

Next generation C&L band equipment includes the day one deployment of both C&L band optimized/ready hardware. These new photonic systems include C&L band amplified spontaneous emission (ASE) that provides “full power” to the entire C&L band spectrum from day one. As working channels are added to the C&L system, the ASE for that spectrum is replaced by the new channel, keeping the total power constant. This means that once the new C&L photonic system has been engineered and turned up, no additional link engineering or optimization is required. The operator can count on stable, predictable performance across the lifetime of the system, regardless of the channel count.

### Integrated Amplified Spontaneous Emission (ASE) – optimized C&L-band upgrade



**Figure 7: Integrated Amplified Spontaneous Emission (ASE) – Optimized C&L-Band Upgrade**

From a deployment perspective, all intermediate sites (line amplifier and dynamic gain equipment (DGE)) would include full C&L band capable equipment day one. The “terminals” would include the C-Band WSS and C-Band amplifiers as well as both C&L band ASE hardware. Using this configuration, the C-Band can be fully populated with transponders from day one while the L-Band (with ASE) is idle. Then, when capacity demands require upgrade to L-Band, only the terminal sites are visited where the L-Band WSS and amplifier hardware is added. Since the intermediate line amplifier and DGE sites are deployed

full C&L band day one, these sites do not require any additional visits or work. All the limitations of the first-generation L-Band have been addressed and corrected with these new C&L band photonic systems.

An additional benefit of having all the C&L line amplifiers and DGEs already deployed is L-Band upgrade velocity. Since the L-Band upgrades only require a visit to the terminal sites, these upgrades can happen in a very rapid manner.

#### **4.4. Additional Fiber Diagnostics**

Today's new photonic systems also include "onboard" fiber diagnostic equipment. Every new photonic element (terminals, line amps, and DGEs) now includes an integrated optical time domain reflectometer (OTDR) that can measure fiber length, fiber loss, reflective events, and chromatic dispersion.

Additionally, with the support of the built in ASE, the non-linear properties of the fiber can also be evaluated in real time. These enhanced fiber characterization capabilities will make photonic layer turn ups more accurate and provide the NMS with in-service, real time fiber information permitting a host of improvements in network operations and maximizing network capacity.

### **5. Software and Automation**

Software and automation continue to increase its role in our lives, from an alarm clock on our smart phones to automated control of our home thermostat. Likewise, software and automation are becoming more critical to the efficient operations of network operators as well. The following are next generation uses of software and automation that will become essential to operators' ability to work efficiently.

#### **5.1. Software – From Planning to Operating**

Telecom operators have always struggled with software. There were different software platforms for every vendors' equipment. There was an Element Management System (EMS) for each different network element. There were Network Management Systems (NMS) for every platform. These different systems typically did not communicate with each other and as a result, interoperability between different vendors' EMS/NMS systems and the equipment under those systems was difficult to affect. These systems were basically good at one thing . . . managing their specific equipment and the signals going through them.

Today's next generation NMS systems take a much more holistic approach. These new NMSs are designed with "open" in mind. North and south bound interfaces allow operators to utilize single system orchestration across multi-vendor domains. Additionally, these systems are no longer just for operations, alarms, management, and provisioning. Next generation network management systems have been architected to additionally provide support for the entire network life cycle from planning to deployment to operations to optimization.

#### **5.2. Planning Equals Deployment**

In the past, planning was considered a necessary evil. There were never any good planning tools. So, everyone "had a spreadsheet". When design and planning are completed offline, the probability of the plan and the network becoming out of sync is very high.

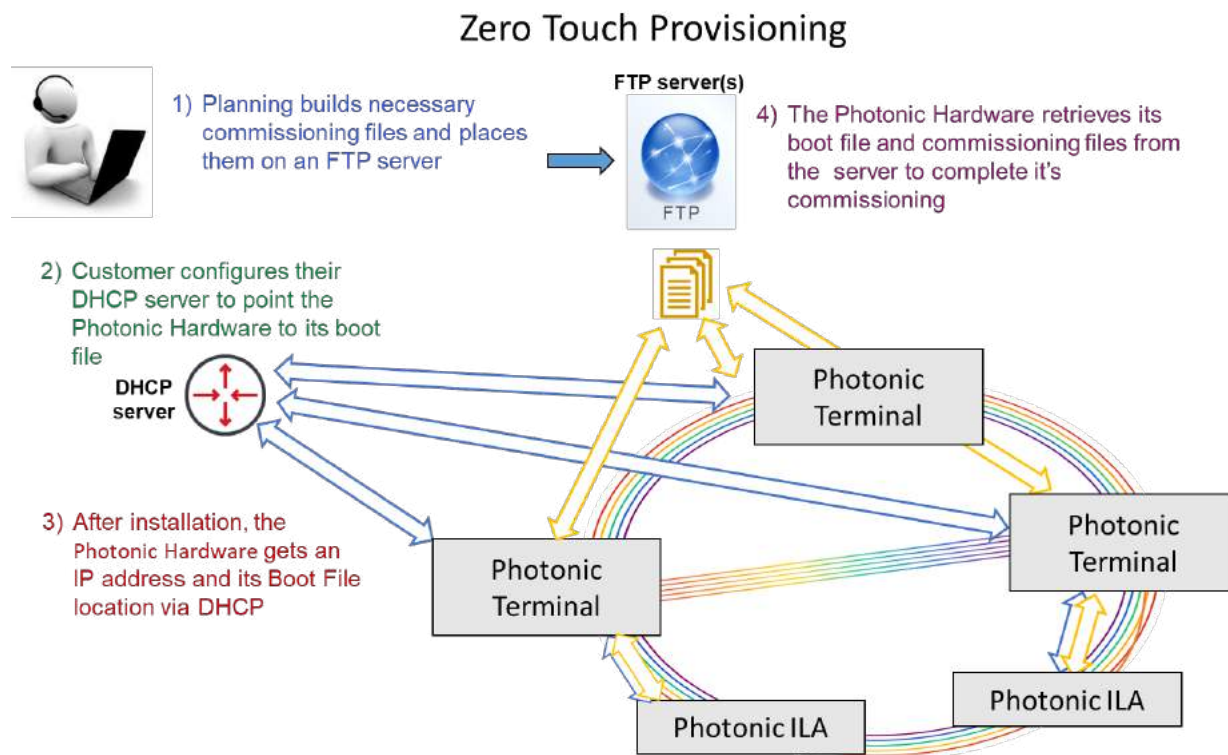
The next generation NMS has integrated the planning function into the NMS. Now, design and planning can be done on the NMS itself. Those designs and plans are stored on the NMS which assures that there is complete agreement between the designs and the network. When new planning is needed, the network is the starting point for the new plan and all previously created designs are considered in the new plan. As plans are implemented and become deployed in the network, the NMS is aware of these adds. The

ability of the NMS to be the gatekeeper is especially helpful in companies with large planning groups. The synchronization between planned and deployed keeps everyone on the same page thus reducing waste from duplicate, overlapping, or incompatible designs.

### 5.3. Zero Touch Provisioning

Once a new design has been created, how is that design processed into deployed and operational equipment? In the past, the design was probably transferred to a spreadsheet and a drawing was made which was then handed off to the deployment team who re-typed the provisioning data from the spreadsheet into the new equipment. This method was inefficient and potentially inaccurate with many manual operations required.

In today's environment where NMS planning has been completed, the NMS planned design is converted into a group of Zero Touch Provisioning (ZTP) files by that same NMS. The NMS sends the necessary ZTP files to a designated FTP server. After physical installation and power-up of the equipment is complete (including the connection of fiber), the installer provides the equipment with an IP address and the location of the FTP server. The new network element retrieves its boot file and commissioning file and then self-installs the commissioning information. Once the ZTP process is complete, the element reboots and is fully functional and ready for use.



**Figure 8: Zero Touch Provisioning**

Another aspect of photonic ZTP is the ability of the amplifiers to self-characterize the fiber to which they are connected. The onboard OTDR provides advanced Fiber Characterization (FC) to thoroughly characterize the fiber plant connected to each amplifier. This FC information is provided back to the NMS. Using the configuration tools built into the NMS, optimized amplifier provisioning information can be calculated and provided back to the photonic hardware on a span-by-span basis to achieve the best

possible performance. This software based, photonic optimization results in the highest possible spectral efficiency for that fiber path.

Finally, post turn up, these same fiber characterization tools allow operators to monitor their fiber plant to preemptively respond to issues and provide detailed information to isolate those issues when they arise.

Since the ZTP information delivered to each network element has been validated by the NMS, it is highly accurate. The accuracy of this information will reduce the overall time required to turn up new network elements. When bandwidth demands require operators to deploy a full photonic system every year, ZTP will be essential to accomplishing this deployment in an accurate, timely manner.

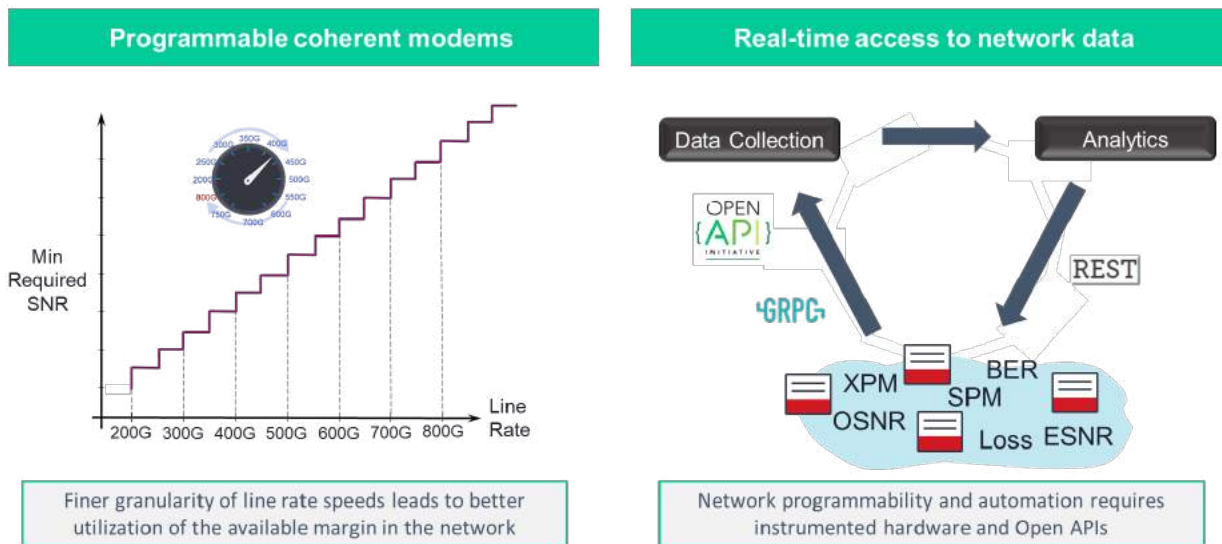
#### **5.4. Margin Mining**

On older fixed grid photonic systems, it was difficult to determine available transponder Rx OSNR margin. In some cases, error corrections could be counted and compared to pre-FEC numbers. Then, using complicated math, operators were able to confirm that their transponder channels were running within acceptable margins. However, even when these calculations were available, there was limited ability to change the characteristics or performance of those channels.

Today, advance software features coupled with next generation photonics and newer baud and bit-rate adjustable transponders provide many optimization options to network operators. One of these newer software implementations is margin mining.

Next generation photonics can provide enhanced operational information including fiber loss, dispersion, fiber types, span counts, BER, and SNR margin. Additionally, next generation transponders can change both the baud and bit rates to maximize performance. Using these parameters, the NMS can now identify available margin and quantify that margin for use in the network. Margin mining can be used to identify this available margin and increase a transponder's bit rate improving the throughput on the link. Margin mining could also be used to reduce the guard band between channels while using the current baud and bit rate. In both cases, there is an increase to spectral efficiency which will result in longer system longevity.

## Optimizing capacity for SNR requires programmable optics AND real-time analytics



**Figure 9: Optimizing capacity for SNR requires programmable optics AND real-time analytics**

## 6. Conclusion

There is a limit to how many bits can be sent down an optical fiber. This limit, known as the ‘Shannon Limit’ is defined as the maximum rate at which data can be sent over a medium with zero errors. Technologies such as coherent optics allow operators to get closer to this theoretical limit. However, moving forward, the gains in spectral efficiency that will be achieved with future generations of coherent technology are diminishing. This will require alternative approaches and ideas to deal with network scalability challenges. Using updated designs, next generation hardware, and software tools, network operators will be able to extend the life of their networks as well as deploy new networks more quickly, efficiently, and accurately. Using these ideas, operators can start reducing the slope of their spectral usage curve while at the same time, deploying new photonic networks with higher efficiency.

# Abbreviations

|        |                                                      |
|--------|------------------------------------------------------|
| A/D    | add/drop                                             |
| BER    | bit error rate                                       |
| CAGR   | compound annual growth rate                          |
| C-Band | conventional fiber optic band 1530 nm to 1565 nm     |
| Demux  | demultiplexing                                       |
| DGE    | dynamic gain (flattening) equipment                  |
| EMS    | element management system                            |
| FC     | fiber characterization                               |
| FEC    | forward error correction                             |
| FTP    | file transfer protocol                               |
| GBaud  | gigabaud                                             |
| Gbps   | gigabits per second                                  |
| Ghz    | gigahertz                                            |
| ITU    | International Telecommunications Union               |
| ILA    | intermediate line amplifier                          |
| L-Band | long or extended fiber optic band 1565 nm to 1625 nm |
| MSO    | multiple systems operator                            |
| Mux    | multiplexing                                         |
| O-E-O  | optical-electrical-optical                           |
| NMS    | network management system                            |
| OSNR   | optical signal to noise ratio                        |
| Rx     | receive                                              |
| SNR    | signal to noise ratio                                |
| SRS    | stimulated raman scattering                          |
| Tbps   | terabits per second                                  |
| WSS    | wavelength selectable switch                         |
| ZTP    | zero touch provisioning                              |

# Bibliography & References

<sup>1</sup> Cisco's VNI Complete Forecast Highlights: Global – 2022 Forecast Highlights; Cisco, 2018; [https://www.cisco.com/c/dam/m/en\\_us/solutions/service-provider/vni-forecast-highlights/pdf/Global\\_2022\\_Forecast\\_Highlights.pdf](https://www.cisco.com/c/dam/m/en_us/solutions/service-provider/vni-forecast-highlights/pdf/Global_2022_Forecast_Highlights.pdf)

# **IoT & Cognitive Computing Approach to Managing Equipment**

**Connect → Predict → Repair → Optimize**

A Technical Paper prepared for SCTE•ISBE by

**Tom Woginrich**

Associate Partner, Global Center of Competence for Asset Optimization  
International Business Machines (IBM)

# Table of Contents

| <b>Title</b>                                                     | <b>Page Number</b> |
|------------------------------------------------------------------|--------------------|
| 1. Introduction – Connected Devices .....                        | 3                  |
| 1.1. Connected Devices, Analytics, and Cognitive Computing ..... | 3                  |
| 1.2. Connected Devices and People .....                          | 5                  |
| 2. Maintaining Equipment.....                                    | 6                  |
| 2.1. Connect .....                                               | 7                  |
| 2.2. Predict .....                                               | 7                  |
| 2.3. Repair.....                                                 | 8                  |
| 2.4. Optimize .....                                              | 10                 |
| 3. High level Architecture .....                                 | 11                 |
| 3.1. Technical Architecture.....                                 | 12                 |
| Abbreviations .....                                              | 12                 |

## List of Figures

| <b>Title</b>                                                                                                                                               | <b>Page Number</b> |
|------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| Figure 1 – Business Benefits will be realized with a combination of traditional analytics and cognitive approaches.....                                    | 4                  |
| Figure 2 – Connected Devices + Analytics + Cognitive Computing + People brings Business Value .....                                                        | 5                  |
| Figure 3 – The data needed to holistically make optimal decisions is widely disbursed, full of inconsistencies and historically difficult to combine ..... | 6                  |
| Figure 4 – Connect → Predict → Repair → Optimize.....                                                                                                      | 7                  |
| Figure 4 – Organization of knowledge sources.....                                                                                                          | 9                  |
| Figure 5 – Operationalization of AI technology through natural language queries (NLQ), troubleshooting diagnosis and augmented reality (AR) .....          | 10                 |
| Figure 7 – High-level architecture .....                                                                                                                   | 12                 |



# 1. Introduction – Connected Devices

The Internet of Things (IoT) is radically changing the way businesses operate and people interact. Billions of devices, sensors, and chips—many of them simple, everyday objects—can communicate with us and each other. Hospitals can monitor and regulate pacemakers at long distance, factories can automatically address production line issues, and hotels can adjust temperature and lighting based on a guest's preferences.

The IoT is changing the nature of products and equipment, as well as customers' expectations of their partners and suppliers. Consumers are pressing for greater accountability and better outcomes from manufacturers for the products and services they provide. As more devices connect the user experience directly back to the manufacturer, it creates expectations that the manufacturer is aware a problem is occurring, and potentially even before it occurs. Organizations like John Deere used to “simply” provide equipment to farmers, yet these farmers are now expecting John Deere to help them use the equipment more effectively. These rapidly changing expectations create both opportunities and challenges. Michael Porter and James Heppelmann summarized it well in *The Harvard Business Review*:

*Once composed solely of mechanical and electrical parts, products have become complex systems that combine hardware, sensors, data storage, microprocessors, software, and connectivity in myriad ways. These “smart, connected products”—made possible by vast improvements in processing power and device miniaturization and by the network benefits of ubiquitous wireless connectivity—have unleashed a new era of competition. (How Smart, Connected Products are Transforming Competition, 2014)*

## 1.1. Connected Devices, Analytics, and Cognitive Computing

IoT is not just about connecting devices. It is a critical first step, but ultimately an enabling one for what follows. An IoT of billions of devices and sensors is pointless unless it effectively utilizes the flow of information. Most of the data generated from these devices will never be used as it simply cannot be understood in time to act. As described in a recent McKinsey Global Institute report:

*Currently, most IoT data are not used. For example, on an oil rig that has 30,000 sensors, only 1 percent of the data are examined. That's because this information is used mostly to detect and control anomalies—not for optimization and prediction, which provide the greatest value. (The Value of Digitizing the Physical World, 2015)*

IoT is about the creation of new insights by analyzing the data that the connected devices generate. First movers are already shifting focus to make sense of the data as it is generated and leverage it for higher value functions. Companies are adopting analytics solutions to process the data for diagnosing problems, predicting outages before they occur, and prescribing solutions.

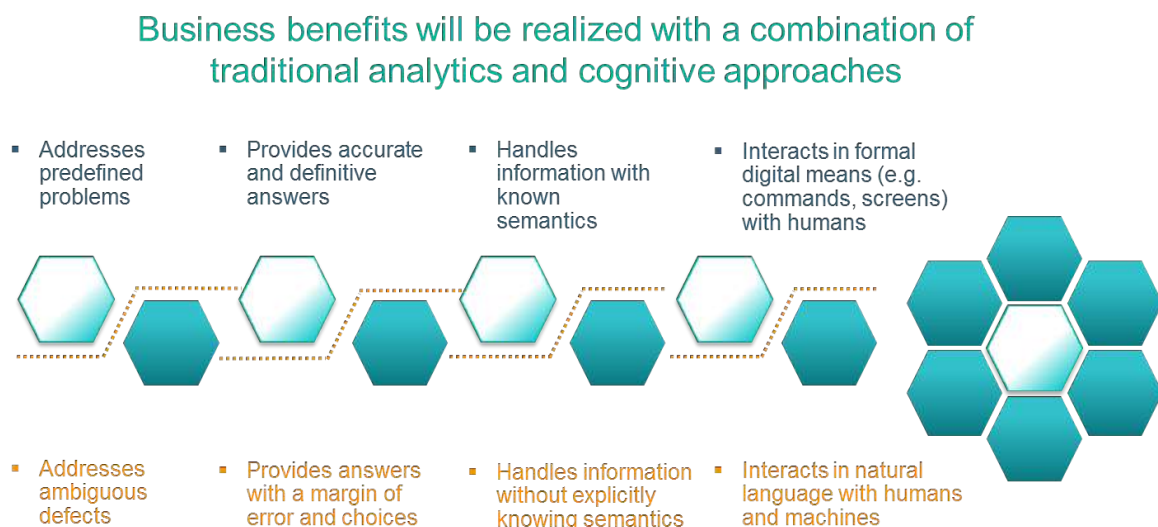
This use of analytics is also having an impact on existing sensor placements as well their modus operandi – including new sensors or relocating sensors for the purpose of predictive analytics. Classic use of

sensors have been to ensure proper operation and quality control. Connected device capabilities now include lifecycle reliability characteristics monitoring – prevent catastrophic failures and extend the useful life thus reducing total cost of operations (TCO).

However, traditional analytics have limitations when it comes to identifying relationships among unstructured data, such as equipment manuals, work orders, and service history. This is where cognitive computing comes into play.

Cognitive augmented intelligence infuses thinking into objects, systems, and process. It provides new opportunities to ingest and analyze data, including unstructured data. Through cognitive capabilities, systems can continuously learn, interpret, and respond. Advanced cognitive systems can communicate via natural language... see, listen and learn.

Cognitive computing and traditional analytics are two sides of the same coin and both are required to achieve business results:



**Figure 1 – Business Benefits will be realized with a combination of traditional analytics and cognitive approaches**

The combination of IoT with cognitive capabilities opens new avenues for clients to boost operational performance, enhance customer experience, and lead industry transformation:

- Work smarter by providing employees cognitive insights from a company's collective knowledge of information
- Give equipment and devices the power to reason and learn, and the ability to interact naturally with people
- Facilitate expedient repair of broken or faulty equipment via a cognitive assistant


## 1.2. Connected Devices and People

Capitalizing on the IoT will further accelerate the shift to what we call Digital Operations. Digital Operations can be defined as a deliberate strategy of transforming businesses through analytics-led, Internet of Things enabled, business and operating models focused on predictive decision making and optimized efficiency.

Few organizations, however, believe they are ready for that scale of change. IBM's Institute for Business Value (IBV) has documented C-level concerns that the pace and breadth of change are beyond their current people's ability to address them. Capitalizing on the IoT will require organizations to rethink how their process engineering integrates IoT services into manufacturing and delivery. Additional data science and quantitative analysis skills, which are already in short supply, will be required to make smart use of the pending flood of IoT generated information streams. Further, IBV research determined that organizations are not sufficiently building the people resources required to address changes resulting from IoT.

Any IoT solution must consider the implications on the organization's people or it will ultimately be limiting and ineffective. Both data and people are essential linchpins in a holistic IoT solution because of the requirement for the right information, processed the right way, driving the right actions by machines and humans. Simply put – the best data in the world will not help unless people can put it to effective use. Since cost and talent issues prevent “throwing people at it”, the effectiveness of the employee needs to be raised to collectively become a smarter organization. We are not implying the replacement of humans by machines, but cognitive systems communicating insights that expand the boundaries of what can be known and acted on. IoT will continue to change the nature of how machines and humans work together.

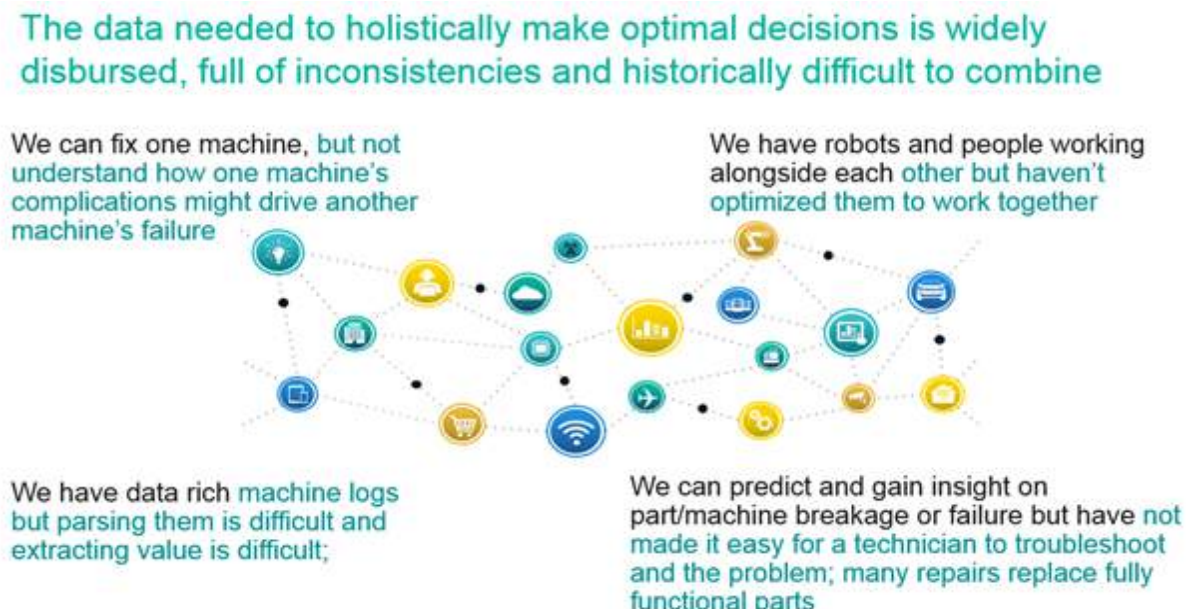
By its very nature, the IoT is a connected web of interdependencies. An effective IoT strategy requires addressing both data and human challenges to succeed. Aggregating these elements brings true business value.



**Figure 2 – Connected Devices + Analytics + Cognitive Computing + People brings Business Value**

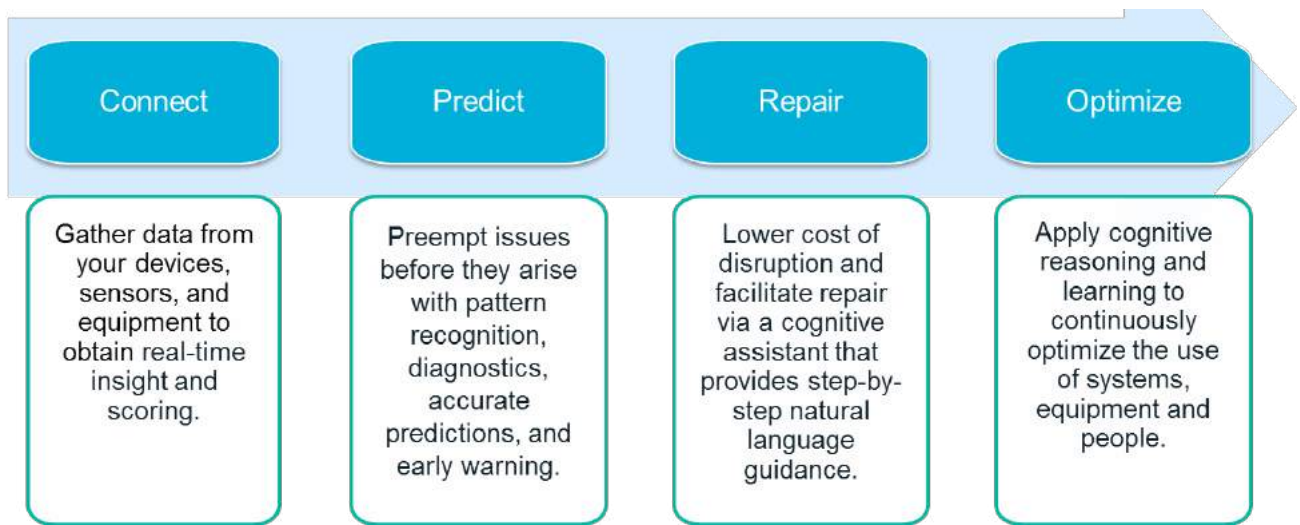
## 2. Maintaining Equipment

By nature, the equipment management ecosystem has had its own current set of challenges before the IoT. As the interdependencies and complexity of operations has increased (whether in an Industrial, Commercial, Mobile, or Consumer-based setting) over the last 2+ decades with advancing control technologies, it has created a new set of challenges:



**Figure 3 – The data needed to holistically make optimal decisions is widely disbursed, full of inconsistencies and historically difficult to combine**

**connect** to equipment, **predict** equipment or system failures before they occur, facilitate **repair**, and **optimize** their operations.



**Figure 4 – Connect → Predict → Repair → Optimize**

## **2.1. Connect**

For multiple decades, electronic sensors have been a staple in the manufacturing industry, in transportation, in consumer products, in utilities and across all other industries. Availability and accessibility to that data has been done in the form of many disparate systems, databases and solutions. The IoT platform enables quick connection of devices, sensors and data and infuses intelligence into applications and services. The Predict and Optimize pillars are big consumers of both realtime and historic data, and the intravenous access provided by the IoT platform is transforming equipment management.

The IoT platform leverages AI to enable real-time insights, natural language processing, machine learning, and video/image/audio analytics. The AI takes the data inputs from connected devices and other sources to uncover patterns and insights previously unattainable. It recommends and implements actions such as safety controls, energy management, and quality assurance.

Key elements include communications with multiple protocols, analytics tools, device data storage, application development environment, data protection, and integration with third party devices and platforms. This approach utilizes the data received from connected devices to feed real-time data into the Predict pillar.

## **2.2. Predict**

Suggesting the obvious, identifying and solving problems before they occur is less expensive and prevents operational disruption.

Work with multiple data sources to generate models targeted at predicting asset failure or quality issues so your organization can avoid costly downtime and reduce maintenance costs. Predictive analytics can detect even minor anomalies and failure patterns to determine the assets and operational processes that are at the greatest risk of problems or failure. This early identification of potential concerns helps you deploy limited resources more cost effectively, maximize equipment uptime and enhance quality and supply chain processes, ultimately improving customer satisfaction.

This approach utilizes analytics to:

- Provide predictive indicators of pending asset failures
- Quickly identify primary variables as part of root cause analysis
- Minimize product quality and reliability issues via early warning indicators
- Optimize spare-parts inventory and help to normalize an unpredictable supply chain
- Provide advance indicators of warranty claims to increase customer satisfaction
- Enhance sales and operations planning to reduce operations costs

It is the output generated from predictive analytics that becomes the input to the Repair pillar to initiate human action as needed.

### **2.3. Repair**

As the IoT transforms customer expectations and competitive pressures, demands on the people working within those industries also increases. This further stresses the human side of product and service delivery and creates problems moving to a digital operations model.

At a field engagement level, the flood of data creates additional challenges for effective service delivery. Field Technicians are often overwhelmed with information from a variety of internal and external data sources.

## Internal Knowledge



## Field Notes



## Collaboration Sources



## Customer Data



**Figure 5 – Organization of knowledge sources**

As machinery and equipment evolves in the IoT era and becomes more complex, the requisite skills and knowledge to repair the systems also increases. Often, there is a lack of effective methods to diagnose and resolve issues quickly out in the field. This impacts company operations and customer satisfaction and potentially drives poor use of parts and resources.

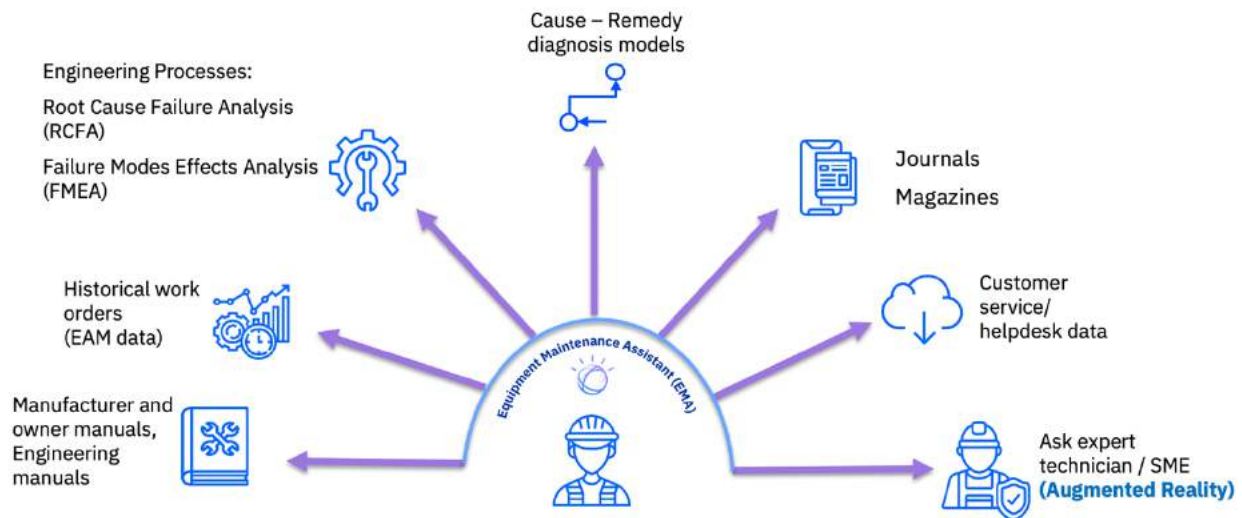
To solve these challenges, companies are leveraging AI as a cognitive advisor for field service technicians. AI ingests various forms of information related to the equipment and repair of it; service manuals, knowledge base, service history, support information, forums, client information, etc. AI learns from this information to develop the best of course of action for repairing the equipment. When an outage or breakage occurs, the technicians can engage AI directly over a simple interface including diagnosis and chat, connecting to current internal systems for specific data. The technician can communicate with AI in natural language to obtain help in diagnosing the problem and step-by-step repair instructions. AI can continue to support the field technician if there are questions or additional information needed as they work on the repair. Field technicians can interact with AI conventionally via a keyboard, or as useful, verbally via Speech-to-Text. AI understands the context that the field technician is asking and can down-select the likely materials and interactions from the Connect and Predict parts of the process that initiated the Repair cycle.

Based on the flexible nature of an IoT Platform building blocks, the solution can be adapted to meet specific business needs. For example, AI could use visual recognition to help identify the type of equipment and the nature of the problem to further expedite repair.

This approach utilizes AI to:

- Ensure problems are fixed correctly the first time – First time to fix metric
- Decrease repair costs by lowering time spent on repair – Mean time to repair (MTTR)

- Optimize training by providing a cognitive tool for teaching – helping to address the aging workforce challenge faced by many industries
- Drive incremental revenue by creating new services opportunities
- Improve customer satisfaction



**Figure 6 – Operationalization of AI technology through natural language queries (NLQ), troubleshooting diagnosis and augmented reality (AR)**

## 2.4. Optimize

This is defined as the Optimize step as an approach that systematically improves outcomes through continual closed-loop learning. Closed-loop learning means that the recommendations are compared to actual outcomes after the actions have been performed. Learning from predictions, while not as widespread as it should be, is not in itself new but this approach addresses a key gap in the process. While this approach learns consistently with machine learning approaches, it also actively learns from the human activity side of the IoT. More specifically, this approach learns from how humans interact with the predictive analytics predictions and how the work was successfully accomplished to both improve the statistical modeling to predict the future and how subsequent engagements provide better context to the people actually doing the work.

Optimize applies cognitive computing and analytics to provide closed-loop learning based on outcomes to constantly improve the usage of systems, equipment, and people. It requires a consultative approach to determine specific areas in the process that could be optimized. There is no “one-size-fits-all” approach and the technology is based on the use case and desired outcome.



Following are some of the areas that we are currently working with clients to optimize their use of machinery and equipment:

- Identify correlations in separate defects among interdependent clusters of systems to improve repair efficiency
- Enhance first-time-fix effectiveness by predicting other components that may fail due to the repair
- Improve cognitive learning based on the outcomes of previous repair jobs for a selected asset or similar repair jobs on like assets
- Leverage machine learning and AI's visual recognition to evaluate manufacturing performance and make recommendations to alter parameters to improve the process

### **3. High level Architecture**

As described above, this approach consists of four pillars; Connect, Predict, Repair, Optimize. The modular approach provides opportunity for a phased approach. However, the value from the aggregation and integration of the parts exceeds the sum of the individual components. Following is a sample architecture connecting these four areas into one solution.

### 3.1. Technical Architecture

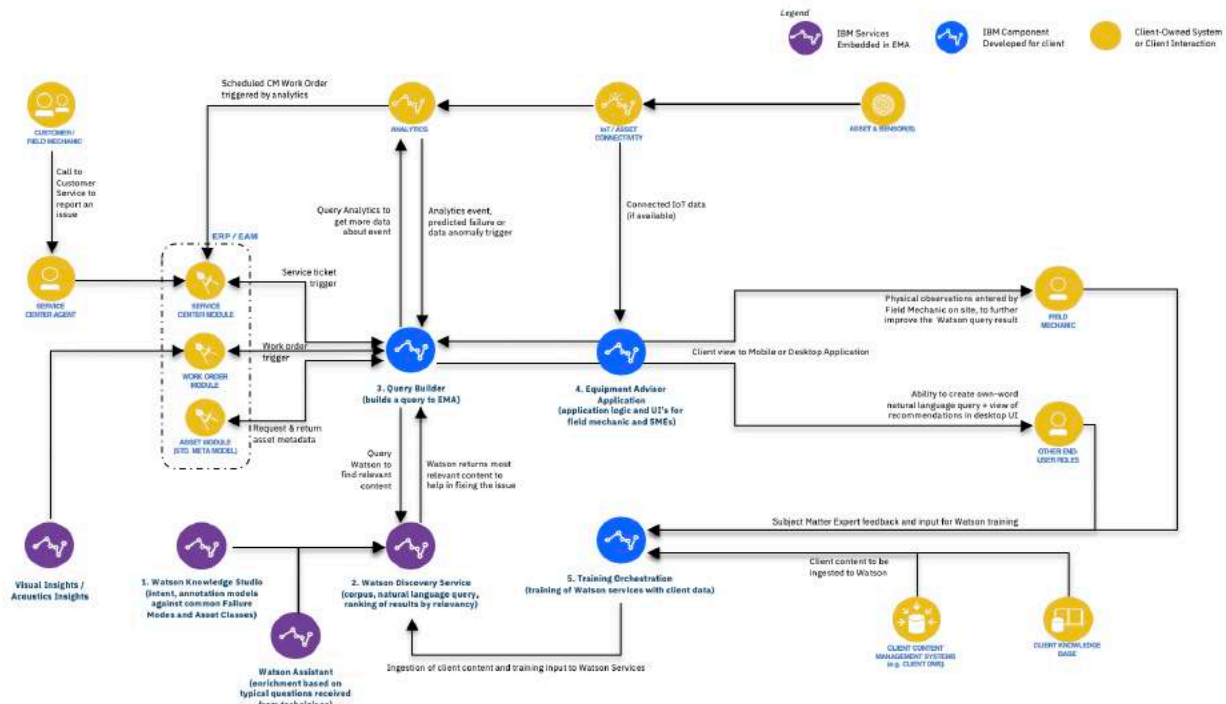


Figure 7 – High-level architecture

In this example, sensors connected to the equipment provide real-time updates to monitor the activity. Predictive analytics are used to evaluate the operational health of the equipment and provide early warnings of trouble. Machine learning continuously sweeps for non-obvious patterns. AI can ingest information from structured and unstructured sources to more quickly identify the issue. Visual cues are provided to display early warnings and if necessary, a repair request is initiated. The field technician leverages AI to obtain step-by-step instructions on the repair using natural language processing and conversation. AI continues to learn and develop a corpus of information through the process which provides continuous improvement.

## Abbreviations

|      |                                               |
|------|-----------------------------------------------|
| AI   | Augmented Intelligence                        |
| IoT  | Internet of Things                            |
| MTTR | Mean Time to Repair                           |
| FTTF | First Time to Fix                             |
| TCO  | Total Cost of Operations                      |
| NLQ  | Natural Language Query                        |
| AR   | Augmented Reality                             |
| SCTE | Society of Cable Telecommunications Engineers |

# **Employing Neural Networks For Improved Root Cause Analysis In Service Provider Clouds**

## **Pros And Cons, Use Cases, General Approach, And Best Practices**

A Technical Paper prepared for SCTE•ISBE by

**Helen Zeng, Ph.D.**

Staff Consulting Solution Architect

VMware

Palo Alto

hzeng@vmware.com

**Robert McIntyre**

Product Development

VMware

Atlanta, Georgia

bmcintyre@vmware.com

# Table of Contents

| Title                                                                      | Page Number |
|----------------------------------------------------------------------------|-------------|
| 1. Abstract .....                                                          | 4           |
| 2. Introduction .....                                                      | 4           |
| 3. Background.....                                                         | 4           |
| 3.1. NetOps, PNM and RCA .....                                             | 4           |
| 3.2. Service Provider Clouds .....                                         | 5           |
| 3.3. Neural Networking .....                                               | 5           |
| 4. Are Neural Networks a Good Fit? .....                                   | 6           |
| 4.1. Modern Network Characteristics .....                                  | 6           |
| 4.1.1. Complex .....                                                       | 6           |
| 4.1.2. Dynamic.....                                                        | 6           |
| 4.1.3. Data-rich .....                                                     | 7           |
| 4.1.4. Real-time .....                                                     | 7           |
| 4.2. No Free Lunch.....                                                    | 7           |
| 5. Existing Frameworks .....                                               | 8           |
| 6. Use Case Scenarios .....                                                | 8           |
| 6.1. Anomaly Detection, RCA .....                                          | 8           |
| 6.2. Performance Management .....                                          | 9           |
| 6.3. Fault Management and Configuration Management .....                   | 9           |
| 6.4. Predictive Maintenance .....                                          | 9           |
| 6.5. Security and Fraud .....                                              | 9           |
| 7. General Neural Networking Approach .....                                | 10          |
| 7.1. Data Collection .....                                                 | 10          |
| 7.2. Data Processing Pipeline.....                                         | 10          |
| 7.3. Further Adjustment .....                                              | 11          |
| 7.4. Data Modeling .....                                                   | 11          |
| 7.5. Continuous Learning.....                                              | 12          |
| 7.6. Model Management.....                                                 | 13          |
| 7.7. Model Optimization and Automation.....                                | 13          |
| 8. Other Limitations.....                                                  | 14          |
| 9. Best Practices for Adoption of Neural Networking .....                  | 14          |
| 9.1. Start Where You Are.....                                              | 14          |
| 9.2. Layer on New Capabilities, Stitched Together at a Workflow-level..... | 14          |
| 9.3. Collaborate, Internally and Externally .....                          | 14          |
| 9.4. Be Metrics-driven: Use Testing to Show Impact over Time .....         | 14          |
| 9.5. Avail Yourself of Expertise When You Need It.....                     | 15          |
| Abbreviations.....                                                         | 15          |
| Bibliography & References .....                                            | 16          |

## List of Figures

| Title                                                             | Page Number |
|-------------------------------------------------------------------|-------------|
| Figure 1 – Multiple Screens Manual Processes.....                 | 5           |
| Figure 2 – Continuous ML Flow .....                               | 8           |
| Figure 3 – Data Analytics Model .....                             | 10          |
| Figure 4 – An Example of Streaming Processing Data Pipeline ..... | 11          |

|                                            |    |
|--------------------------------------------|----|
| Figure 5 – Data Model Development.....     | 12 |
| Figure 6 – Continuous Training Cycle ..... | 12 |

## 1. Abstract

Academics, strategy pundits and entrepreneurs like to talk about the transformative power of artificial intelligence (AI) and the need to accelerate implementations. (AI has been called the “new electricity.”[1]) But service provider operations teams tend to tell a different and more nuanced story. While no one disputes the long-term potential of these technologies, what’s becoming clear is that the complexity of advanced AI, such as neural networking, is difficult to fit into their current paradigms.

The pragmatic way to adopt these technologies is to define a clear use case, start with what you already have, and layer on these technologies in such a way that operations teams can augment existing architectures. A phased approach allows room to focus on learning and enhancing rather than replacing existing network operations (NetOps) practices, including root cause analysis, in order to evaluate costs and benefits. Key to this process is focusing on a better-together approach to provide evidence-based demonstration of value. In this paper, we will introduce neural networking, explain why it is a good fit for today’s networks, and provide use cases from our experience in deploying these technologies in 5G and cable environments. It is more than possible to experiment with this technology, while evolving, rather than reinventing existing skillsets and processes.

## 2. Introduction

A method of problem solving that moves beyond reactive to proactive management, root cause analysis (RCA) has long been a goal of effective NetOps. As cable networks have become increasingly complex, with service delivery stitched across access networks, optical fabric, IP/MPLS backbones and, increasingly, cloud infrastructure, NetOps has evolved accordingly. In a NetOps 2.0 world, RCA has not only become an expected feature, it has needed to become smarter and more effective.

Implementing a proactive and predictive RCA in complex service provider clouds that span networks comprising more than a million devices is no trivial challenge. How do you determine the network radius of “problem spaces” that are compounded in extent by multi-tenancy, network traffic paths, and physical and logical L2/L3 connections? Neural networking-based algorithmic approaches, using non-deterministic and probabilistic models and automated computation, are one promising approach to this challenge. Employed to multi-tiered problem spaces with temporal and spatial aspects, this brand of AI technology is highly effective at executing RCA on top of non-deterministic anomaly detection, all within context.

The prerequisites are network and service discovery, along with a platform for data collection, streaming and processing. When properly set up and implemented over physical and virtual infrastructure, neural networking can drive improved RCA and its positive business effects across multiple use cases. On the other hand, if inadequately understood or lacking in data, neural networks can exceed their limits or result in less control. As with any AI or machine learning (ML)-driven initiative, upfront knowledge is key.

## 3. Background

### 3.1. NetOps, PNM and RCA

There are many aspects to NetOps, from the physical Network Operations Center (NOC) itself to the personnel who work there to the policies employed to the technologies used in managing, monitoring and control any number of discrete or interrelated networks. A common term used in the cable industry used for optimizing NetOps is proactive network maintenance (PNM). As a part of a special initiative covering wired, Wi-Fi and optical technologies, CableLabs has encouraged careful thinking about PNM,

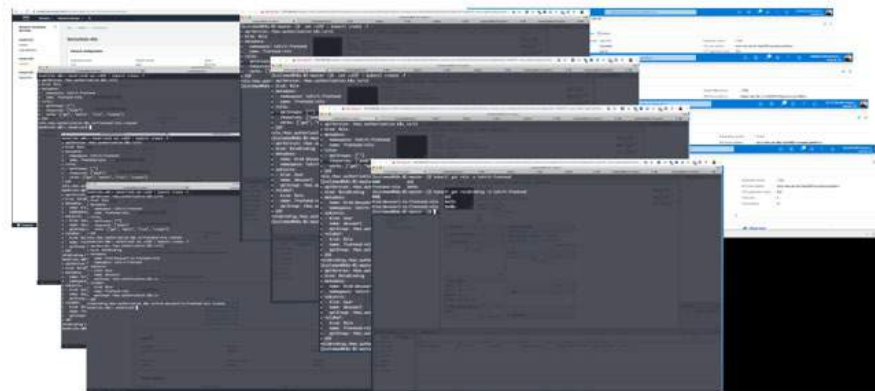
emphasizing the benefits of proactive management.[2] The point of any such exercise, of course, is a higher level of customer service.[3]

The RCA approach to problem solving falls into that broad category of proactive measures, in that it seeks to uncover the underlying cause of faults or incidents affecting network performance; and because the failure to address that cause is likely results in recurring issues. By isolating the fundamental fault associated with a multiplicity of issues, RCA leads to a reduction in immediate number of alarms, as well as a lower number over time. RCA does not by itself remediate problems, but feeds into a process of corrective and preventative action.

### 3.2. Service Provider Clouds

A typical domain for applying the tools of PNM is the cable access plant, for instance DOCSIS-related problems that occur on that part of the network spanning the CMTS or CCAP device and the related modem at the customer premise. Yet with virtualization, that network now extends beyond the physical plant into cloud infrastructure.

The distributed access architecture (DAA) initiative remains one of the industry's most prominent use cases related to virtualization. Being transformed from purpose-built hardware into software enables a cable operator to run the CCAP in a data center on a private cloud. Other possible cases for MSO virtualization include VOD, network PVR, 5G and Multi-access Edge Computing (MEC).[4] These cases increasingly cross industry lines, many applications all requiring high-performance, low-latency networks. The expansion of these networks has placed a tax on traditional NetOps. MSOs have expressed frustration with monolithic management tools, and delight in those that deliver cross-domain results.[5] An extensive physical and cloud environment makes it difficult to rely on siloed legacy tools with manual processes and scripts. (See Figure 1.)



**Figure 1 – Multiple Screens Manual Processes**

### 3.3. Neural Networking

In that context, it is fair to ask whether networks can survive without AI. The real question today, however, is what kind of AI or ML serves best. Neural networking is worth considering. It may appear to be a recent addition to this category, but its genesis goes back decades. Building upon theories of biological neural network, i.e. the brain, the seminal paper on artificial neural networks appeared in

1943.[7] By the late 1990s, engineers were looking at possible applications in manufacturing.[8] Driven in part by innovative big-data ingest and storage capabilities, neural networking in recent years has achieved tremendous progress in areas such as natural language processing and image recognition.

In the cable and telecommunications arenas, other kinds of AI and ML have taken the lead. One example is a trial that involved “operational analytics (OA) and machine intelligence (MI)” to predict service impairments in near real-time, which used Support Vector Machine (SVM) classification, spectral clustering, tree-based taxonomy and related tools.[9] Like these uses of advanced data analytics, neural networks can very well approximate the non-deterministic, stochastic nature of various network scenarios.

Neural networking is typically supervised, i.e. it requires setting up and training a data model. A key characteristic is that these networks involve layers of input, hidden and output neural nodes, each connected to others, with certain weighed coefficients. Besides feed-forward networks, there are other kinds of deep-learning models, including recursive neural networks (RNN) and convolutional neural networks (CNN), etc. But neural networking can be unsupervised as well, especially when no labeled data is available. The system tries to find patterns and form clusters in a meaningful way. It involves mapping the continuous random variables to discrete representations, while neural networks are capable of achieving that through their ability to converge.

## **4. Are Neural Networks a Good Fit?**

### **4.1. Modern Network Characteristics**

There are a number of reasons for service providers of all stripes to consider the applicability neural networking for key NetOps problems, such as RCA. One overriding reason is that for many, networks are simply not what they were only a few years ago. Nor have they reached an endpoint. Monitoring and analytic tools and platforms simply need to keep pace with advancing network technologies. that now increasingly bear these characteristics:

#### **4.1.1. Complex**

A conventional approach to network monitoring and management has involved applying simple thresholds to classifying physical layer HFC performance within normal and abnormal parameters. That remains an effective way to trigger alarms, but insufficient when the goal is to isolate underlying problems affecting services that touch heterogeneous networks, each of which may have its own siloed data arranged in unique formats. Assessing and weighing multiple inputs, and becoming even smarter over time, is possible with the computational power and adaptability of a neural network.

#### **4.1.2. Dynamic**

The disaggregation of once-unified equipment, such as the CMTS or CCAP device, into core and remote elements, is an indication of not only growing network complexity but also accelerated change. The addition of remote PHY devices (RPDs) entails new configurations and adjusted topology adjustments, while a virtualized CMTS core now depends upon cloud infrastructure that may rapidly scale up and down. Edge wireless, IoT devices and even CPE contribute more unknowns to the mix. Neural networks are adept at handling dynamic change.



### **4.1.3. Data-rich**

The sheer volume and variety of data emanating from today's networks is one of the biggest drivers in the search for new approaches. Only a few years ago, network data mean SNMP queries that delivered a limited amount of information every fifteen minutes, and of that very little was useful. MSOs had tools that could handle those challenges. The data challenges now facing many MSOs are beyond their scope. Neural networks actually work better with inputs provided by big data analytics.

### **4.1.4. Real-time**

Exacting performance in areas such as latency matter to many end customers, whether the scenario is cell backhaul or internet gaming. Data about those services matters, too. But even if MSOs were able to staff their NOCs with the sharpest analysts, the rate at which performance data arrives and the need to process them with utmost speed is beyond human capabilities. Manual review of Syslogs and support cases do not scale, and status-quo analytic tools fall short. To respond to network behavior real-time, service providers need certain automatic ML-based mechanisms, and neural networking is a very good ML model.

## **4.2. No Free Lunch**

Neural networking is not the only advanced AI and ML-based analytic tool available. For supervised classification problems, there are some commonly used ML algorithms, such as SVM, decision tree, logistic regression, etc. For decision tree, rigidity of the model is the common problem which easily leads to over-fitting. Although they can be trained to be accurate, once given new data, they may jump to wrong conclusions. That could happen through creating tree branches that follow certain order of precedence. Ensemble methods, however, can alleviate this problem.

SVM performs data separation either linearly or non-linearly, which highly relies on the choice of kernel function. It makes a model difficult to scale well to a large dataset. While logistic regression tends to underperform when there are multiple or non-linear decision boundaries, it is not flexible enough to naturally capture more complex relationships. Neural networking is more flexible as it has the capability to approximate almost any scenario accurately by adjusting hidden layers and hidden nodes.

A common criticism of neural networks is that they operate like black boxes. While a decision tree is relatively easy to understand, it is hard to explain how a neural network arrives a particular conclusion. Neural networks do well with large amounts of data, which makes them a good fit for advanced communications networks; but the converse applies, they do less well in scenarios that are data-deprived. Computationally powerful, they can also require more time to train than traditional ML algorithms, which can increase their cost. Likewise, operating them requires some expertise in data science.

Shortchanging knowledge and taking a set-it-and-forget-it approach can be tempting but dangerous. A single-engine propeller aircraft can be operated and maintained manually; whereas a commercial jet airliner requires NetOps 2.0-level systems and procedures. Yet when pilots are unable to override these systems and are unprepared for edge cases that may arise, they cede too much control to the automation that was designed to optimize flight operations.[10]

A similar charge could be leveled against a neural network, which generally speaking is a “non-convex optimization problem,” especially when the activation function is non-linear. (Since weights are permutable across layers there are multiple solutions for any minima.) Sometimes, domain knowledge is needed to make sounding judgements when selecting reasonable ML tuning parameters. The takeaway is that it is important, whether building one of these platforms yourself or working with a partner, to understand the technology's potential benefits, its operating characteristics and its limits.

## 5. Existing Frameworks

How do existing service assurance platforms handle RCA? A crucial prerequisite of any framework is data collection, streaming processing and data modeling. As for RCA more directly, one currently effective approach is to rely on multi-dimensional matrices of deterministic models that create ‘signatures’ related to symptoms and problems. The idea is to correlate events and alarms with known patterns and signatures to identify where the root of the problem lies. Other best practices include continuous self-updating; adapting to dynamic workloads, network configurations and inventory; and integrating with orchestration tools for auto-remediation and incident management for support workflows.

What status quo platforms typically do not yet have is a continuous flow of anomaly detection, leading to RCA for the anomaly, and finally providing prescription for the problem. A non-deterministic, neural networking engine can extend while enhancing other effective techniques for anomaly detection, RCA and problem prescription. The neural networking ML utilizes streaming processing output, although the training can be done off-line. Then streaming processing engine picks up ML result, through multi-stages to generate real-time results in an automated fashion. This paper focuses on RCA, while the problem prescription is out of scope.

Below is an example of flow using neural networking in radio access network (RAN) throughout the whole data analytics process. (See Figure 2.)

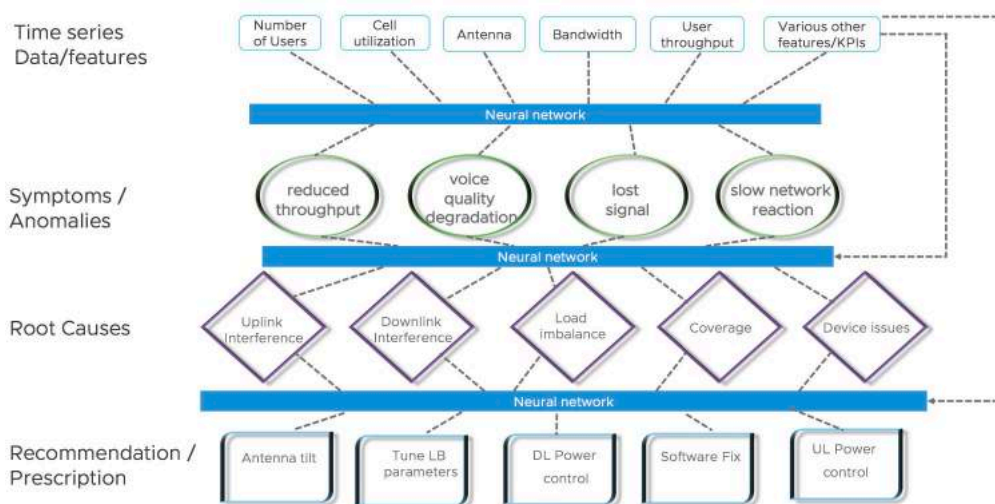


Figure 2 – Continuous ML Flow

## 6. Use Case Scenarios

### 6.1. Anomaly Detection, RCA

Neural networking works well for anomaly detection, which is both a powerful technique and a prime use case. The detection of anomalies, or statistical outliers, is applicable in complex, on-prem/off-prem and dynamic environments where operators are interested in long-term trends, aberrations such as spikes, and cross-correlation with other inputs. In the case of having no global formula to define anomaly, neural networking is able to learn the network performance down to the subscriber level over a long-enough period of time and determine the anomaly. All of that information can help drive effective RCA, which

could also be considered a use case. Anomaly detection can be supervised, unsupervised or semi-supervised, depending upon the use of a training data set. Like RCA, it figures in several other use-case scenarios.

## **6.2. Performance Management**

To track the delivery of any number of services across a large subscriber base, a service provider today needs to receive concurrent data feeds from a massive number of network elements, correlate those with user session data, and then calculate real-time key performance indicators (KPIs). The big data platform enables an operator to achieve those tasks. Neural networking is built on top of the platform to combine the KPIs with application-specific inputs and policies to deliver network visibility, anomaly detection, and real-time predictive network intelligence.

## **6.3. Fault Management and Configuration Management**

A fault management system must have a comprehensive picture of the network topology, which is related to configuration management. There are thousands of ways a network can fail, go sideways or skip a beat. Hardware failure, connectivity loss, and power outages are some types of network faults. Let's consider one, the misconfiguration of a network device, such as an RPD. Anomaly detection might first indicate that subscriber data throughout is not matching up with other session-level parameters. A system equipped with neural networks for anomaly detection might detect conditions where average user throughput dropped below a certain level, while the channel utilization was abnormally high, and DOCSIS sub-carrier utilization low. Identifying the condition and the underlying cause then enables a recommendation for device reconfiguration.

Another example is a cellular tower fault. Each cell has neighbors, and if one cell is in trouble, then the user traffic is distributed to the neighbor cells, at which point the neighbor cells will show an abnormal increase of traffic amount and maybe congestion. Some network faults directly impact or even block service delivery. Neural networking can certainly help network fault management to detect/predict fault and further diagnose the source and type of the fault. The subsequent action is to automatically trigger fault correction, or to at least prevent incidents from happening in the future.

## **6.4. Predictive Maintenance**

In its early forms, neural networking was seen as a way to help monitor manufacturing processes, predict failures and schedule maintenance on aging equipment. It likewise applies to telecommunication and IT network equipment, which can experience transient or permanent hardware or software malfunction. One concern involving new proliferating IoT endpoints is battery life.[11] Neural networking could assist in managing this and other aspects of IoT, including the detection of missing data transmission and related RCA, either as part of an MSO's own network or a managed service to businesses.

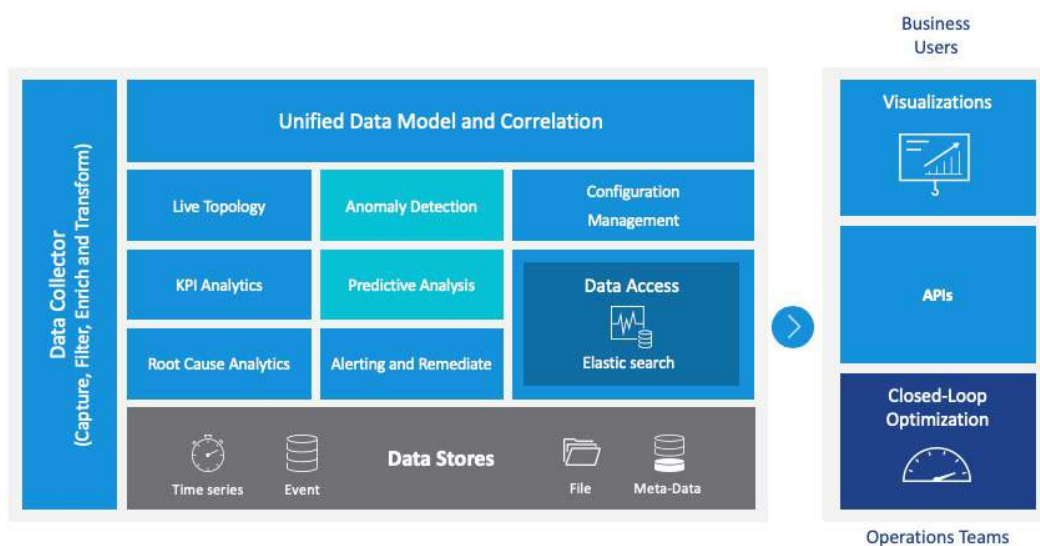
## **6.5. Security and Fraud**

In the area of intrusion detection, network security experts have combined anomaly detection models with various deep neural networking structures.[12] Models could detect strong login behavior, too many DNS requests or frequent changing of temporary IP or ID during a session. Operators themselves are motivated to prevent telecommunications fraud on their own networks and to deliver high-value managed security services to other businesses. To detect fraud, the data covers several domains, such as network usage in multiple types of networks, financial, personal information, location information, etc. Fraudsters are good at frequently changing their behavior to match the 'normal' pattern of the network to circumvent the security rules, which increases the level of difficulty for fraud detection. In addition to detecting

identity theft and fraudulent activity, network operations intelligence has also had to track the location and time of fraud. The data-mining capability of neural networking has made it a candidate for providing better information to MSOs about unsanctioned use of telecom networks, whether for financial gain, abuse of services, ghosting, tampering, cloning or other fraud.[13]

## 7. General Neural Networking Approach

A neural network-based service assurance recognizes that the world has changed. Service-related data once arrived slowly and with limited utility. Then services became IP-centric, and it was economically and technically feasible to ingest and store exponentially more data than before. Applying neural networking for anomaly detection, RCA and other uses entails completing a sequence of tasks: data collection (and network discovery), data processing, data modeling and follow-up adjustments. (See Figure 3.)



**Figure 3 – Data Analytics Model**

### 7.1. Data Collection

This task involves collecting fault and performance metrics for each network component at the user level. There are traces, events as performance metrics. The fields include start and end-time on each of 1000s of pieces of equipment, user ID, traffic type (voice/video/data) IP address and location, throughput, traffic volume, etc. For security and fraud use cases, typically call detail records (CDR) and network management system (NMS) logs are needed. CDR and logs have subscriber-level call details and activity records in a certain period, i.e., once a month. Network devices could use various industry standard data collection methods, such as SNMP, REST, NEP, etc., for streaming data sources. The time stamp is important for model training and prediction. Data collectors may span multiple vendors, with their own implementations and formats, as well across network silos. What generally appears in the NetConf standard are configuration files, counters and information about the state of the device.

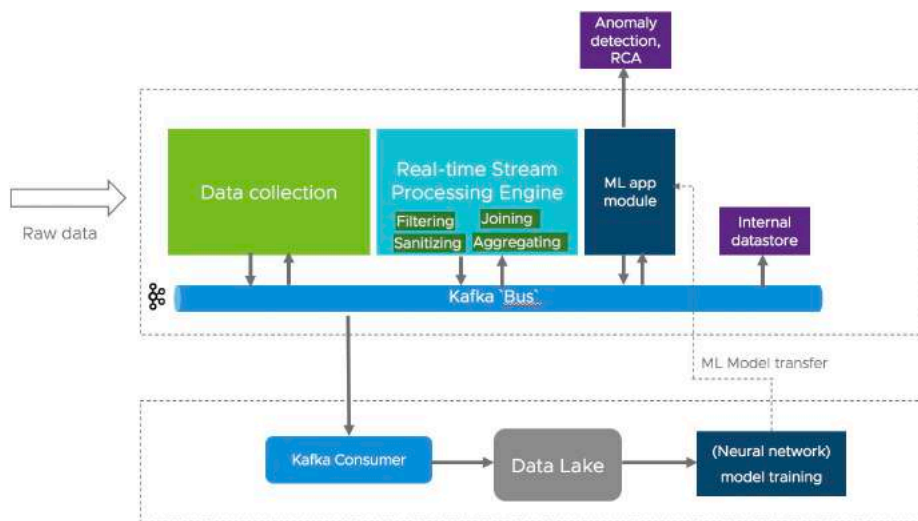
### 7.2. Data Processing Pipeline

After setting up the collectors, the next step is to set up real-time streaming and processing of the time-stamped data into a data pipeline and connecting to a database. The process of data joining can be a

challenge in the streaming domain, with SQL joins being especially slow and difficult. Yet standard big-data and schema-agnostic database formats have helped smooth the preparation phase. Related tasks include filtering, sanitizing the data, as well as running network topology discovery to associate data with device location. There is a Kafka bus to connect various big-data components in the pipeline. Those components exchange data/information through the Kafka bus. Some big data frameworks are Apache Spark, Twitter Heron, Apache Flink, etc.

A common way of storing the processed data is to ingest data via Kafka into a data lake, which will be used for neural networking model training. The streaming processing engine is equipped with a ML module. The module can get the ML model result and uses streaming data to generate prediction results, such as anomaly detection or RCA.

Below is one example of streaming data processing pipeline including neural networking model training engine. (See Figure 4.) Note that data lake is not the only way to store the data; there are other ways which do not require a data lake, such as tiered storage in Kafka. Also, some neural networking ML use cases do not require the results to be fed back to real-time streaming processing engine; they can certainly run separately.



**Figure 4 – An Example of Streaming Processing Data Pipeline**

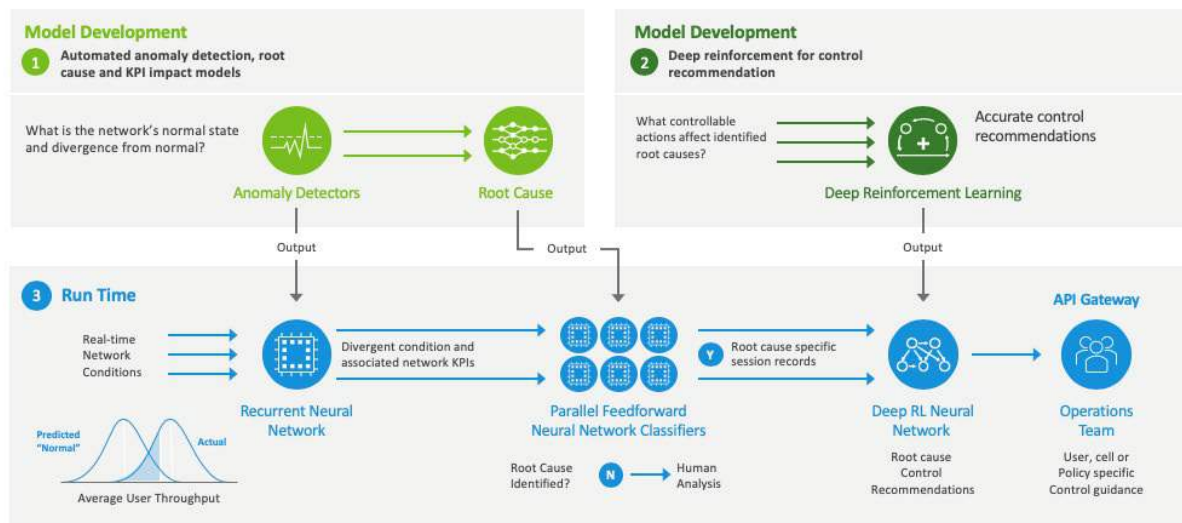
### 7.3. Further Adjustment

Other data management techniques may be required. During data preparation, in the case of missing data, average or default data may be inserted, or the fields could be ignored. Several approaches can be taken in the case of data skewness, such as oversampling, or synthetic minority oversampling technique (SMOTE) to compensate for an imbalance within the set. Neural networks tend to make better predictions when trained on balanced data.

### 7.4. Data Modeling

Correlating chosen KPIs (containing time stamp, user ID, traffic type, etc.) with different network components is important to establishing an end-to-end view. A critical part of building a model is identifying features, which correspond to parameters or counters in events. These could be throughput, signal strength, transmit power, or whichever ones best align with the use case in question. All of those


features are associated with a single time stamp that represents training data that are fed into the model. Data scientists can optimize the model’s “fit” by adjusting the weights and layers of the middle nodes, which may have distinctive properties of their own, such as the ability to do “convolutional” math or look back across the network in a “recursive” manner. If so engineered, the model will deliver outcomes that determine a root cause or detect anomalies. (See Figure 5.) There are various ML framework libraries to draw upon, including TensorFlow, Theano, Deeplearning4j, Keras, sklearn, etc.



**Figure 5 – Data Model Development**

## 7.5. Continuous Learning

The model training is a continuous process. We should proactively compare ML predicted result with ground truth and then adjust the model if necessary. A separate effort can be placed on continuously enhancing the model, and then feedback the new model in the subsequent data analytics flow. (See Figure 6.)



**Figure 6 – Continuous Training Cycle**



## 7.6. Model Management

After a neural networking ML model is delivered, the following criteria need to be evaluated:

- ML performance, which can be evaluated by comparing the prediction result with ground truth and quantified by mathematical formula.
- What features are used? Can the feature list be changed?
- Are the training data generalized enough?
- How much training data is enough?
- Is there any over-fitting or under-fitting?
- What are the optimal hyper parameters? Are they generalized or specific to particular scenarios?
- Is human input needed? For example, sometimes human help is needed for training data labeling, or domain knowledge is needed to customize the ML prediction result.

The above considerations should be taken into the model management. Further actions like model versioning, customization, refinement, profile, performance recording, dependency should be incorporated when deploying models:

- Model versioning: If a model is only suitable for old data, a version is used to avoid using the wrong model on new data.
- Model customization: ML models are normally consistent across deployments but can be customized to take into account market-specific conditions. For example, different hyper-parameter lists, or values are used for different markets (due to different RF bands, geographical situation or other factors); different feature lists or weights for different markets or clusters, etc. Another case is that network performance should follow some theoretical rule, but sometimes an AI/ML-driven prediction lands outside of the reasonable boundary; in which case human customization or adjustment is needed.
- Model profile: Certain mechanisms to prevent over-fitting, under-fitting, and ML software framework can be recorded for evaluating ML model performance.
- Model performance: This helps determine how good a model is. It can be evaluated in various ways, such as, F1-score, area under the curve (AUC), receiver operating characteristics (ROC) curve, learning curve, R2-score for regression algorithm, etc. Also the learning curve can help determine how much training-set data are needed or are sufficient.
- Model dependency: When new or novel patterns are detected, the associated event stream can be used as labeled data by domain experts (which suggests human input) for re-training of models using supervised learning. This human-input dependency can be incorporated into model-training software.

## 7.7. Model Optimization and Automation

Machine learning helps realize a big portion of network operations automation. But the ML model itself can also be optimized and automated. In the case of neural networking, the model performance heavily relies upon the hyper parameters, including the number of hidden layers, the number of nodes in each hidden layer, the learning rate, etc. In a production environment, it is crucial to automate the process of selecting optimum tuning parameters. Bayesian optimization appears to be an efficient choice, as it not only helps find the vector of hyper parameters that result in a neural network with the lowest error, but also reduces considerably the time spent on model tuning. There are other optimization tools or ML libraries available, as well.

## 8. Other Limitations

As mentioned earlier, while neural networking thrives on large amounts of data, in some cases there may not be enough. Challenges surrounding data joining have also been noted. Mistakes on a key or value could result in the time stamp being off, which would make it difficult to correlate data belonging to the same user session. One solution is to set a slightly wider time range. Finally, there are always going to be false positives and false negatives. The tradeoff between precision and recall cannot be avoided. In setting up the model and use case, operators should take note of the use case and consider any false positives.

## 9. Best Practices for Adoption of Neural Networking

Experience across the wider service provider industry has revealed an effective way to approach this powerful and challenging technology. If you are motivated to successfully adopt neural networking for RCA, consider the following key principles:

### 9.1. Start Where You Are

Allow yourself room to experiment. Consider a data-driven before/after business case that can be used to showcase progress and share learnings. We recommend starting with one of the use cases mentioned above, perhaps anomaly detection, being applicable in many scenarios.

### 9.2. Layer on New Capabilities, Stitched Together at a Workflow-level

Early adopters are thrilled to discover that neural networking can deliver insights that were previously impossible to obtain. Yet the difficulty in explaining how it works can hinder practical and actionable impact. By framing the technology as an augmentation, on a second screen or a dashboard with all anomalies detected by neural networking tagged as an overlay, operations teams can get started with the technology in a way that stages risks and provides ample room for learning and evolving, both in terms of testing and tuning the models, but also to mitigate false positives while the models are being developed.

### 9.3. Collaborate, Internally and Externally

Given the rate at which these technologies are evolving, and the centrality of model definition and evolution, it pays to collaborate with other teams both internally and externally. In our examples here, many of the models applicable to any network scenario were developed initially for 5G interference scenarios but are equally applicable to DOCSIS environments and LLX deployments. Look for ways to collaborate across silos and share insights.

### 9.4. Be Metrics-driven: Use Testing to Show Impact over Time

Neural networking is not a reinvention as much as an evolution, and in to demonstrate value and prioritize investments in these technologies, vanguard teams who are seeing the most success with implementations have taken “show me” approaches to proving value. The most straightforward way is to define a clear use case and employ A/B testing to demonstrate the value with before-and-after metrics. Prioritize future steps accordingly.



## 9.5. Avail Yourself of Expertise When You Need It

Because of the learning curve and pitfalls of developing, training, evolving, and explaining neural networking models, having access to the right level of expertise is crucial in getting these programs off the ground. It can be time-consuming to identify domain experts who are adept at translating the data science principles into the context of RCA. Plan ahead.

# Abbreviations

|         |                                           |
|---------|-------------------------------------------|
| 5G      | fifth-generation cellular wireless        |
| AI      | artificial intelligence                   |
| AUC     | area under the curve                      |
| CCAP    | converged cable access platform           |
| CDR     | call detail records                       |
| CMTS    | cable modem termination system            |
| CNN     | convolutional neural network              |
| CPE     | customer premises equipment               |
| DNS     | domain name system                        |
| HFC     | hybrid/fiber coax                         |
| IoT     | internet of things                        |
| IP      | internet protocol                         |
| KPI     | key performance indicator                 |
| LLX     | low latency Xhaul                         |
| MEC     | multi-access edge compute                 |
| MI      | machine intelligence                      |
| ML      | machine learning                          |
| MPLS    | multi-protocol label switching            |
| MSO     | multiple systems operator                 |
| NETCONF | network configuration protocol            |
| NMS     | network management system                 |
| NetOps  | network operations                        |
| NOC     | network operations center                 |
| OA      | operations analysis                       |
| PNM     | proactive network maintenance             |
| PVR     | personal video recorder                   |
| RAN     | radio access network                      |
| RCA     | root-cause analysis                       |
| REST    | representational state transfer           |
| ROC     | receiver operating characteristics        |
| RNN     | recursive neural network                  |
| RPD     | remote PHY device                         |
| SNMP    | simple network management protocol        |
| SMOTE   | synthetic minority oversampling technique |
| SQL     | structured query language                 |
| SVM     | support-vector machine                    |
| VOD     | video on demand                           |

# Bibliography & References

- [1] Shana Lynch, "Andrew Ng: Why AI is the New Electricity," Insights, Stanford Business, March 11, 2017
- [2] Jason Rupe, "A General-Purpose Operations Cost Model to Support Proactive Network Maintenance and More," SCTE-ISBE, 2019
- [3] Andrew J. Milley, "Proactive Customer Maintenance," SCTE-ISBE, 2019
- [4] Andrew Bender, "A Roadmap for Virtualization in HFC Networks," SCTE-ISBE, 2019
- [5] Jeff Baumgartner, "Comcast's AI action extends to the network core," Light Reading, July 20, 2020
- [6] Claudio Righetti, et al., "Can Future Networks Survive Without Artificial Intelligence," SCTE-ISBE, 2019
- [7] Warren McCulloch and Walter Pitts, "A logical calculus of the ideas immanent in nervous activity," The bulletin of mathematical biophysics, vol. 5, pp. 115-133, 1943
- [8] Tiago A. Piedras Lopes, Antonio Carlos R. Troyman, "Neural Networks on Predictive Maintenance of Turbomachinery," IFAC Proceedings Volumes, Vol. 30, Issue 18, August 1997.
- [9] Justin Watson and Roger Brooks, "Predicting Service Impairments from Set-top Box Errors in Near Real-Time and What to Do About It," SCTE-ISBE, 2018
- [10] Nadeem, "The Deadly Price of the Automation Paradox," The Walrus, September 26, 2019
- [11] Joe Rodolico, "Methods to Maximize IoT Battery Life," SCTE-ISBE, 2019
- [12] Naseer, et al., "Enhanced Network Anomaly Detection Based on Deep Neural Network," IEEE Xplore, Aug 17, 2018.
- [13] Gurunadham, "Identifying Telecommunication Deception using Neural Networks through Data Mining," International Journal of Engineering and Techniques, Vol. 3, Issue 6, Dec 2017.

# **City Of Dublin: Lessons From A Smart City Private Network Deployment**

A Technical Paper prepared for SCTE•ISBE by

**Ladan Pickering**

Professional Services Technical Lead  
Fujitsu Network Communications  
2801 Telecom Pkwy, Richardson, TX 75082  
972-479-2371  
Ladan.pickering@fujitsu.com

# Table of Contents

| <b>Title</b>                                                | <b>Page Number</b> |
|-------------------------------------------------------------|--------------------|
| 1. Introduction.....                                        | 3                  |
| 2. Smart City Market in North America.....                  | 3                  |
| 3. Why is Digital Transformation Important to Cities? ..... | 4                  |
| 4. Parking Occupancy App.....                               | 4                  |
| 5. Dublin Architecture .....                                | 5                  |
| 6. Why CBRS? .....                                          | 6                  |
| 7. Lesson #1 – Deployment .....                             | 7                  |
| 8. Lesson #2 – It is all about data .....                   | 7                  |
| 9. Lesson #3 – Environmental challenges .....               | 8                  |
| 10. COVID-19 - Compliance .....                             | 8                  |
| 11. Conclusion.....                                         | 9                  |
| Abbreviations .....                                         | 9                  |

## List of Figures

| <b>Title</b>                                   | <b>Page Number</b> |
|------------------------------------------------|--------------------|
| Figure 1- Smart City Market North America..... | 3                  |
| Figure 2- Smart Parking, City of Dublin .....  | 5                  |
| Figure 3- Darby Street Lot.....                | 5                  |
| Figure 4- Dublin Architectre .....             | 6                  |
| Figure 5- CBRS Architecture.....               | 7                  |
| Figure 6 - COVID-19 Compliance .....           | 8                  |

## List of Tables

| <b>Title</b>                | <b>Page Number</b> |
|-----------------------------|--------------------|
| Table 1- 5G NA Market ..... | 4                  |

## 1. Introduction

According to Harbor Research and analysis by Fujitsu, the Smart City market will reach \$5.6B in North America by 2025.


Cities own hundreds of departments that are using software targeted to their functions. These datasets are generally local to the organization. In addition, real-time access to digital representation of city's physical assets are rare.

Using the CBRS-band (Citizens Band Radio Service), cable companies can build private networks for cities to bring together disparate systems and make it economically feasible to digitize and access physical assets in real-time. Visibility to data from multiple systems will provide insight to improve operational efficiencies, provide better experiences for citizens, and improve environmental impact.

In this paper, we will discuss the lessons learned from our private network deployment and Parking Occupancy App that was deployed by City of Dublin. We will also discuss the positive unintended consequences of the private network and App during COVID-19 outbreak.

## 2. Smart City Market in North America

\$ Millions



**Figure 1- Smart City Market North America**

**Table 1- 5G NA Market**

| Smart Cities NA | 2020     | 2021     | 2022      | 2023      | 2024      | 2025       |
|-----------------|----------|----------|-----------|-----------|-----------|------------|
| Public          | \$425.8  | \$782.7  | \$1,285.2 | \$1,987.2 | \$2,954.1 | \$4,270.6  |
| Private         | \$151.48 | \$273.17 | \$438.81  | \$661.97  | \$957.79  | \$1,347.72 |

Source: Harbor Research & Fujitsu Extrapolation

### **3. Why is Digital Transformation Important to Cities?**

Digital representation of cities' assets such as parking spaces, roads, utilities and fuel consumption is important to the cities because it provides citizens with better experiences, reduces the carbon foot print, improves emergency response, and reduces operational cost.

Historical data based on the digital representation of assets and processes help cities better understand their operations, predict the future, and improve decision making and planning processes.

Software virtualization, low cost sensors, cameras, GPUs, and most importantly the shared spectrum have enabled this digital transformation to be economically feasible.

### **4. Parking Occupancy App**

#### *Customer Problem*

Multiple congested parking lots and lack of parking management impact local businesses

#### *Solution*

Private network and edge computing with video surveillance and analytics



**Figure 2- Smart Parking, City of Dublin**

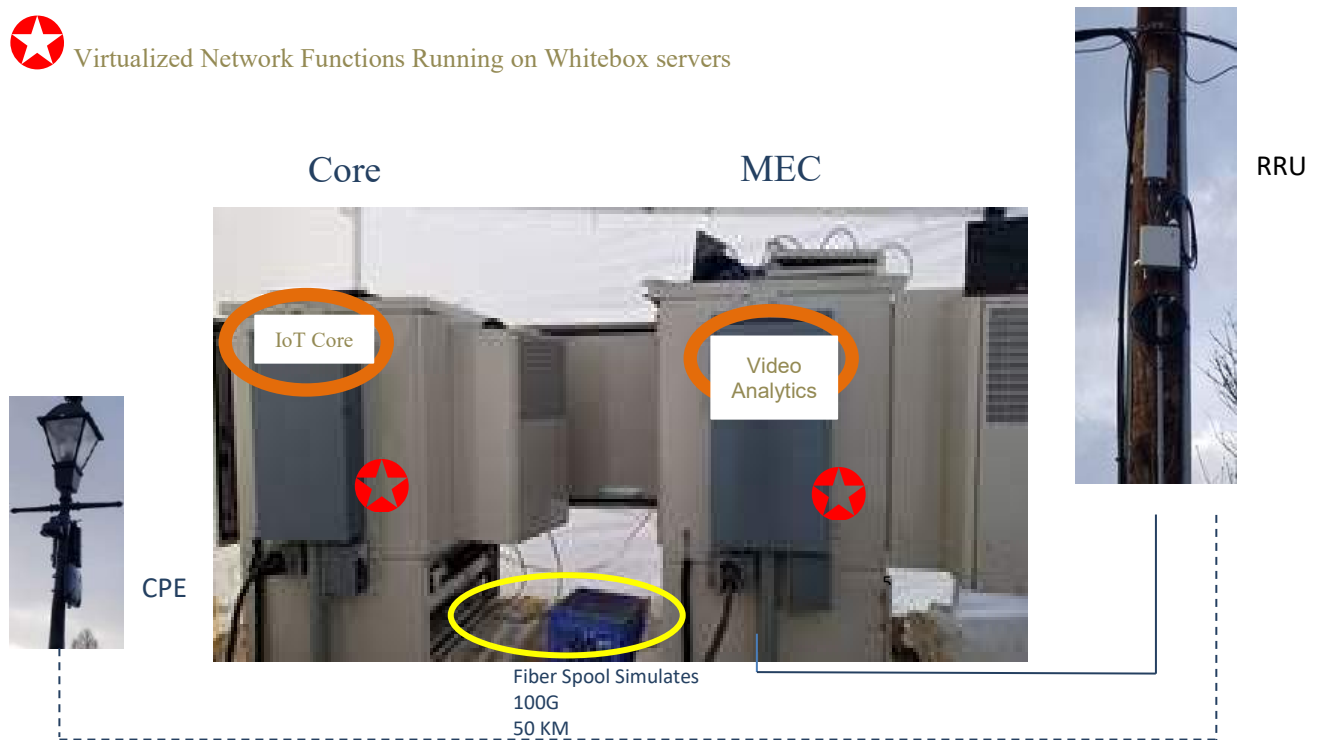


**Figure 3- Darby Street Lot**

## 5. Dublin Architecture

The physical connectivity in the City of Dublin included, Radios, BBU (Base Band Unit), cameras, and CPEs (Customer Premise Equipment). Radios and BBU were installed on a utility Pole at the parking lot. The cameras and CPEs were installed on the city light poles around the parking lot.

Application software used the whilebox servers that provide network functionality. The Video-Analytics software ran on the MEC (Multi-access Edge Compute) node. IoT-framework and application software ran on the Core.



**Figure 4- Dublin Architectre**

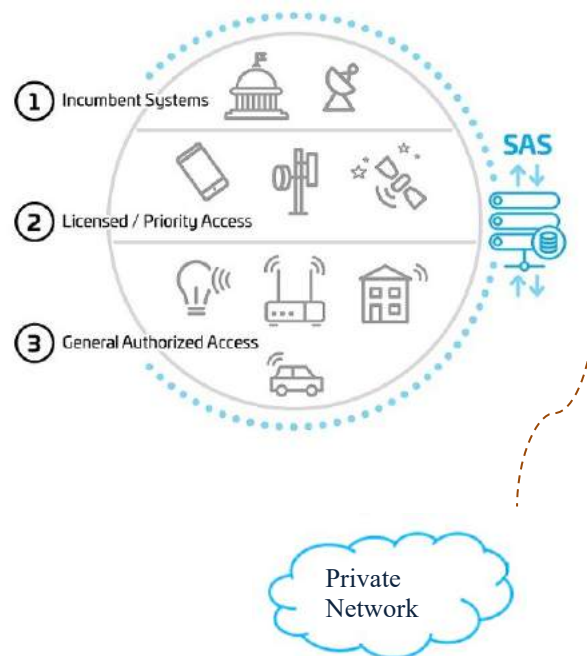
## 6. Why CBRS?

Many critical Enterprise applications that require low-latency, high bandwidth, or a massive number of connections cannot be solved using Wifi (Unlicensed Spectrum) or are economically infeasible to address through public networks (licensed Spectrum).

FCC opened CBRS (Citizen's Broadband Radio Service) which operates at 3.5 GHz band and allows 80 mHz of spectrum to be shared. CBRS is not unlicensed spectrum like Wifi, and it is not licensed spectrum that is controlled by the service providers. It is a new concept where the spectrum is shared in a 3 layer model. When an incumbent is not fully utilizing the spectrum, the Authorize Access Users can jump in and use the spectrum. SAS (Spectrum Access Sharing) is provided through 3<sup>rd</sup> party vendors (Google, Commscope, SONY, ...) that control access.

Cable companies can provide access points that are CBRS-enabled and enter the wireless market to provide Private Networks or Neutral hosts services to enterprises and cities.





**Figure 5- CBRS Architecture**

## **7. Lesson #1 – Deployment**

- Permitting and inspection takes time
- Installation can be expensive
- MEC connectivity should be planned early
- Cameras are not UE's (today)
- Appearance Matters - Customer facing eqpt must look professional and SW must be intuitive

## **8. Lesson #2 – It is all about data**

- Automatic data collection is important
- Site access statistics can alert hacks/security attempts
- Data must be simplified through Graphs/Charts in order to facilitate quick path from “data to decision”
- Open APIs are a must
- Access to real-time images/video streams is important

## 9. Lesson #3 – Environmental challenges

- **Rain & Snow** – Rain drops and snow flakes can attach to the camera and cover portion of the view. This will effect the accuracy of the Video-Analytics
- **Trees & Leaves** – As spring roles around, the leaves on the trees may obscure portion of the camera view.
- **Storms and high wind** – Camera attachments must be highly secure to avoid camera movement during a storm, if the cameras move, adjustment must be made to correct field of view.
- **Sunset/Dusk** – Visibilty of the camera view is effected by the change of lighting during sunset.
- **Length of the attachment poles** – If cameras are attached to short poles, the visibility is reduced and more cameras will be required.

## 10. COVID-19 - Compliance

Although the goal of the project was to better understand the parking congestion, the project allowed City of Dublin staff to remotely monitor compliance to COVID-19 “Stay at Home” orders.

Remote access to the city resources can save lives by eliminating the need for city staff to physically come close to the public.

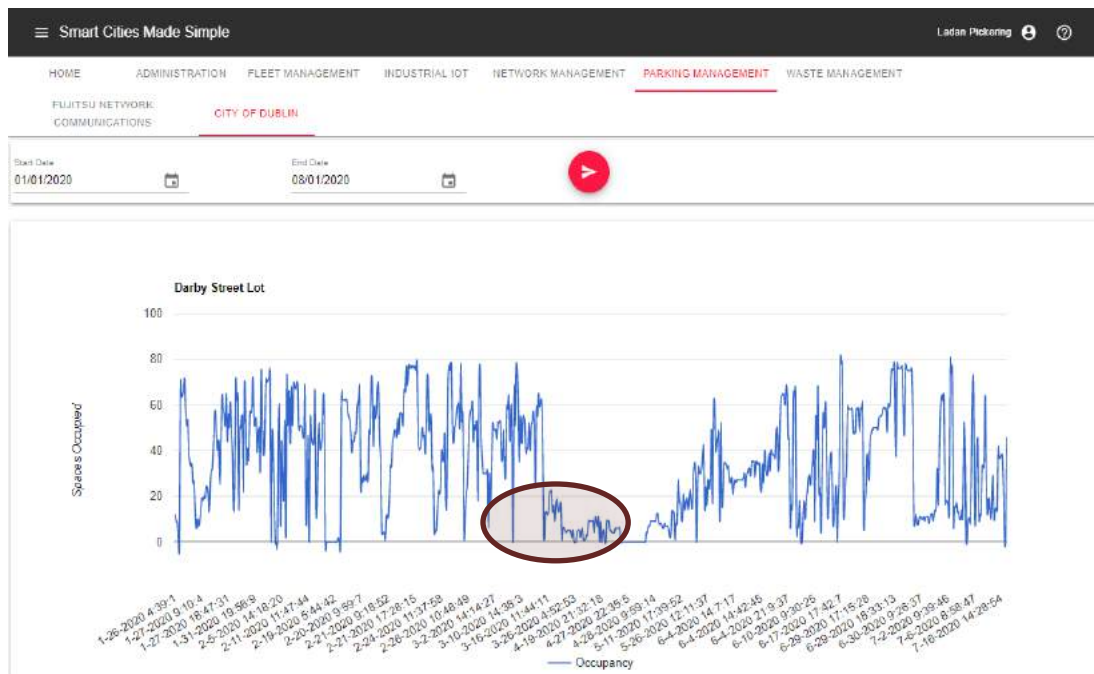


Figure 6 - COVID-19 Compliance

## 11. Conclusion

Private Networks enable deployment of additional Apps. Now that the Dublin Private Network is in place, we can add our COVID-19 based applications (Smart Distancing and Smart Masks) over the air to the network servers to enable additional compliance measurements.

The low cost of CBRS allows cable companies to enter the wireless market by building private networks for enterprises and cities. Once the network is in place, the white box servers that provide network function virtualization can be re-used to run additional application software and therefore, allow cable companies to provide Apps that provide digital twins of Cities' assets.

## Abbreviations

|          |                                               |
|----------|-----------------------------------------------|
| App      | Application Software                          |
| BBU      | Base Band Unit                                |
| COVID-19 | Coronavirus 2019                              |
| CPE      | Customer Premise Equipment                    |
| CBRS     | Citizen's Broadband Radio Service             |
| FCC      | Federal Communications Commission             |
| GPU      | Graphics Processing Unit                      |
| ISBE     | International Society of Broadband Experts    |
| MEC      | Multi-access Edge Computing                   |
| SAS      | Spectrum Access Sharing                       |
| SCTE     | Society of Cable Telecommunications Engineers |
| UE       | User Equipment                                |

# **A Better Platform to Facilitate Remote Patient Monitoring**

A Technical Paper prepared for SCTE•ISBE by

**Jason Page**

Principal Engineer

Charter Communications

6360 Fiddlers Green Circle, Denver, CO 80111

720-699-6236

jason.page@charter.com

## Table of Contents

| <b>Title</b>                                            | <b>Page Number</b> |
|---------------------------------------------------------|--------------------|
| 1. Introduction .....                                   | 3                  |
| 2. What is Remote Patient Monitoring (RPM) .....        | 3                  |
| 3. Why Should Cable Operators Care .....                | 4                  |
| 4. Current Implementations and Their Shortcomings ..... | 4                  |
| 5. A Better Connectivity Solution .....                 | 5                  |
| 6. Why Use a Router .....                               | 5                  |
| 7. Why Use the OpenSync Platform .....                  | 6                  |
| 8. BLE Telemetry .....                                  | 6                  |
| 9. Proof of Concept .....                               | 7                  |
| 10. Areas Requiring Additional Work .....               | 11                 |
| 11. Potential Concerns with the Outlined Approach ..... | 11                 |
| 12. Conclusion .....                                    | 11                 |
| Abbreviations .....                                     | 13                 |

## List of Figures

| <b>Title</b>                                               | <b>Page Number</b> |
|------------------------------------------------------------|--------------------|
| Figure 1 – Device Discovery and Connection .....           | 8                  |
| Figure 2 – Device Connection and Attribute Discovery ..... | 9                  |
| Figure 3 – Data Transfer .....                             | 10                 |

## 1. Introduction

Cable companies are positioned to enable remote monitoring of patient data through connected medical devices. The remote needs of customers due to Coronavirus highlighted the need to offer healthcare services outside of traditional brick and mortar medical facilities. Existing solutions that enable the wireless transfer of data from connected medical devices suffer from numerous shortcomings that limit the adoption of this technology.

The predominant means of wireless connectivity for in-home medical devices is Bluetooth Low Energy. Most manufacturers require the use of a smartphone and a proprietary application to transmit vitals measurements from healthcare devices. This places a heavy burden on a patient's ability to purchase a smartphone and then have the technology savvy to download an application and connect with the healthcare devices. These proprietary applications also lead to data siloes and inconsistent security practices. To overcome this limitation some device manufacturers and service providers resort to using expensive cellular radios and data plans. As a result of these steep barriers many of the most vulnerable patients are unable to participate in the use of this technology.

There is a better way. IoT radios, such as Bluetooth Low Energy (BLE), should be included in traditional Ethernet and Wi-Fi routers. These common access points are a natural bridge between Personal Area Networks (PAN) that are used for most constrained IoT devices' wireless communication and traditional IP based Local Area Networks (LAN) and Wide Area Networks (WAN). This should be combined with a flexible software platform that can provide interfaces for IoT device provisioning, command and control, and telemetry transport.

Utilizing this approach to fulfill remote patient monitoring use cases would enable devices to be onboarded with little or no user intervention. It would also ensure that sensitive healthcare measurements can be sent directly to secure whitelisted endpoints. As administrators of the platform, Cable companies can implement industry standard security practices and provide advanced monitoring and troubleshooting services. This solution removes the technological barriers to entry, improves connection reliability, and can be provided at a much cheaper price point than using cellular.

This paper describes a prototype of a router-based remote patient monitoring system which will be much more effective than the current approaches. It introduces OpenSync, a cloud-agnostic open-source software for the delivery, curation, and management of services for the modern home and the IoT extensions that need to be added to it. It describes the enhancements that must be made to the router and its software stack, as well as what cable companies can do as an industry to enable this new line of business.

## 2. What is Remote Patient Monitoring (RPM)

Remote Patient Monitoring (RPM) is a new technology in the rapidly evolving world of digital healthcare delivery. RPM allows for the collection of patient data outside of traditional brick and mortar healthcare facilities. The ability to regularly take readings of vitals and other physiological measurements enables healthcare professionals to provide custom tailored

treatments and intervene more quickly if patients are not progressing as expected. There are other uses as well: consumers can track their own wellbeing, research and analytic companies can collect anonymized data to improve overall healthcare and design new therapies.

RPM has been enabled by the continual miniaturization of electronics and by the evolution of IoT. Most connected medical devices are battery operated, constrained devices. They typically use Bluetooth Low Energy for wireless communications but sometimes come with cellular or Wi-Fi radios. Most of the devices specialize in taking a single medical measurement and generally attempt to transmit the value as soon as it is available. Examples of RPM devices include connected blood pressure cuffs, pulse oximeters, scales, thermometers, glucometers, and more.

A diverse set of stakeholders is required for RPM to achieve general acceptance and provide value to users. The major stakeholders in the space are patients, doctors, device manufacturers, electronic medical records providers, insurers, patient engagement companies, healthcare exchanges, and many more. Each stakeholder has different interests and responsibilities. These competing factors often lead to data being siloed with a single entity causing patients to get less value from the technology. Coordinating access to medical data and ensuring proper authorization can become very complex. Cable companies can play the coordination role.

### **3. Why Should Cable Operators Care**

Currently there exists no standard for transporting data from connected medical devices to backend systems in the cloud where the data can be used and stored. Cable companies are well positioned to influence the creation of standards and be the connectivity bridge. Cable already has a presence in nearly every home across the country and provides connectivity for millions of connected devices. The next evolution of this connectivity should be to connected medical devices that use personal area networks.

The potential benefits to Cable companies are vast. Cable companies stand to enhance the relationship with their customers and improve their overall reputation by being the guardians to sensitive health information. Internet and advanced wireless connectivity packages can be made stickier since they play a vital role in directly ensuring our customer's wellness. There also exists new lines of revenue through the direct administration of remote patient monitoring services or through partnerships with healthcare systems and providers. This technology also presents an opportunity to engage in the wider smart home industry.

### **4. Current Implementations and Their Shortcomings**

Remote patient monitoring requires the transfer of sensitive data from medical devices to the internet. To facilitate this data transfer two general approaches have been adopted. The first and most common approach is to use a smartphone as a hub and the second approach is to equip medical devices with cellular radios. Each carries its own shortcomings that limit the adoption of this technology.

The smartphone as a hub approach works by utilizing the Bluetooth radio that is ubiquitous in smartphones today. The smartphone acts as a central device and is able to connect to nearby BLE medical devices. Once connected to the medical device the smartphone is able to retrieve readings from the device and send them to the cloud or store them locally. Generally the process of connecting to the device and retrieving measurements is done through a proprietary application provided by the device manufacturer.

The smartphone as a hub approach suffers from a few shortcomings. The first is that it requires users to have a smartphone. Patients are also required to download proprietary applications, turn on Bluetooth, create accounts, adjust permissions and settings, and perform a range of other tasks that may not be simple for many of the most needy. The proprietary nature of the application required to interact with the device often leads to data siloes. Having medical devices from different manufacturers also generally requires separate applications.

To overcome these limitations the second approach of using cellular radios has been adopted. This approach generally places cellular radios directly in connected medical devices. Rather than requiring a hub to enable the transfer of data to the cloud the device can directly transfer the data itself. While this approach eases the technological burden on patients it increases the cost of connected medical devices and requires expensive recurring data plans. It also increases the complexity of the device and leaves the security implementation completely at the discretion of the device manufacturer.

## **5. A Better Connectivity Solution**

An ideal connectivity solution for remote patient monitoring must retain the cost effectiveness of inexpensive Bluetooth radios while minimizing the technical role a user must play. It should provide the patient the ability to authorize access to their medical data but otherwise be transparent to them. The nuance involved in connecting the medical device to a network, controlling the device, and ultimately transmitting readings to the cloud should be automated. The user interface for patients to interact with their medical data should be completely decoupled from the application that transfers data from device to cloud. All device provisioning and association with a particular user should be done by backend systems that require little or no patient involvement.

To have the widest adoption the software should be open source and the router components should be hardware agnostic. An open source solution would allow every connectivity provider and medical device manufacturer to deploy the solution. A hardware agnostic platform would allow components to be cost competitively sourced and limit the leverage of any given vendor. This helps drive down the total bill of materials and avoid procurement obstacles.

## **6. Why Use a Router**

The router is a natural replacement for the use of a smartphone as a hub and its always on nature provides many additional advantages for the collection and transfer of medical data. The router is a relatively inexpensive and common piece of equipment for most households. It currently serves to connect ethernet and Wi-Fi capable devices to the internet. With addition of IoT radios the



router can also be a bridge to connected medical devices and facilitate the transfer of data to the cloud.

Using a router to provide access to the internet for connected medical devices provides an inexpensive connectivity option. This also ensures that existing BLE medical devices can participate. Routers can be centrally and remotely administrated. The always-on access point lends itself nicely to regular collection of connectivity diagnostics. Remote administration can assist in troubleshooting malfunctioning equipment.

## **7. Why Use the OpenSync Platform**

Equipping routers with additional IoT radios is only part of the solution. In addition to the necessary hardware to communicate with connected medical devices a software platform to interact with the devices is also required. Such a platform must be capable of being remotely managed, have a secure pipeline to facilitate telemetry, and be capable of being run on routers. In addition, the software must be deployed to millions of routers and bulletproof.

OpenSync is a software platform that meets the requirements to enable remote patient monitoring. It has baked in support for JSON RPC, a synchronized bi-directional database, and MQTT. These components make the platform extremely flexible, extensible, reliable, and secure. OpenSync and all of its components are open source. It can be run on OpenWRT, RDK, and other Linux based operating systems. OpenSync has also been thoroughly vetted and underpins the routing platforms of a number of major companies.

## **8. BLE Telemetry**

To facilitate the transmission of data from a connected medical device, some means of uniquely identifying the device and some means of communicating with the device are required. In the majority of RPM devices, the means of communication is Bluetooth Low Energy and the device can be uniquely identified by its MAC address. Additional pieces of information that are required for the BLE protocol are the GATT service UUID that is broadcast in advertisement packets, and GATT characteristic UUID(s) that stores the device's value(s) of interest. Additional information is helpful but not necessarily required.

All BLE transactions begin with a device advertising itself. Information included in the advertisement packet is the MAC address of the device and usually a GATT service UUID. The service UUID can effectively be used to determine what type of device is advertising while the MAC address can be used to uniquely identify the exact instance of the device type. Any scanning devices that hear this advertisement can then issue a connect request to the remote device if they are interested in interacting with it. Upon establishment of the connection the central device can then issue commands to the peripheral device. These commands can assist in learning more about the capabilities of the device, reading values from the peripheral, writing values to the peripheral, or enabling notifications.

In a remote patient monitoring use case a typical transaction begins with a device turning on when a user engages it. This could be a user stepping on a scale, pressing a button on a blood pressure cuff, or placing a pulse oximeter on their finger tip. The device begins taking a measurement and when the first reading becomes available the device begins advertising itself. Nearby devices that are listening for these advertisements can issue a connection request in response. When the listening device connects it enables notifications on the medical device and the medical device then transmits values as they become available. This could be a stream of values as is the case with a pulse oximeter or a single value as is the case with a weight scale. When the medical device is done taking its measurement or otherwise disengaged by the user it disconnects from the listening device and goes to sleep. Transactions can grow more complicated when adding in authorization or bonding but this is the basis for the transfer of readings from a connected medical device to an internet connected device.

## 9. Proof of Concept

To prove the viability of RPM on a router and through OpenSync, the Emerging Technology team at Charter Communications designed and implemented a proof of concept. The design used a Raspberry Pi to act as a router and added a Silicon Labs Bluetooth capable radio via USB. The main software component is a custom build of OpenSync with extensions for IoT. These extensions included schema updates, a Bluetooth Low Energy hardware abstraction layer (HAL), a centralized RPM manager, and specialized plugins to assist the manager application. Supporting software components included a RESTful API, MQTT to Kafka connector, WebSocket server and a simple user interface.

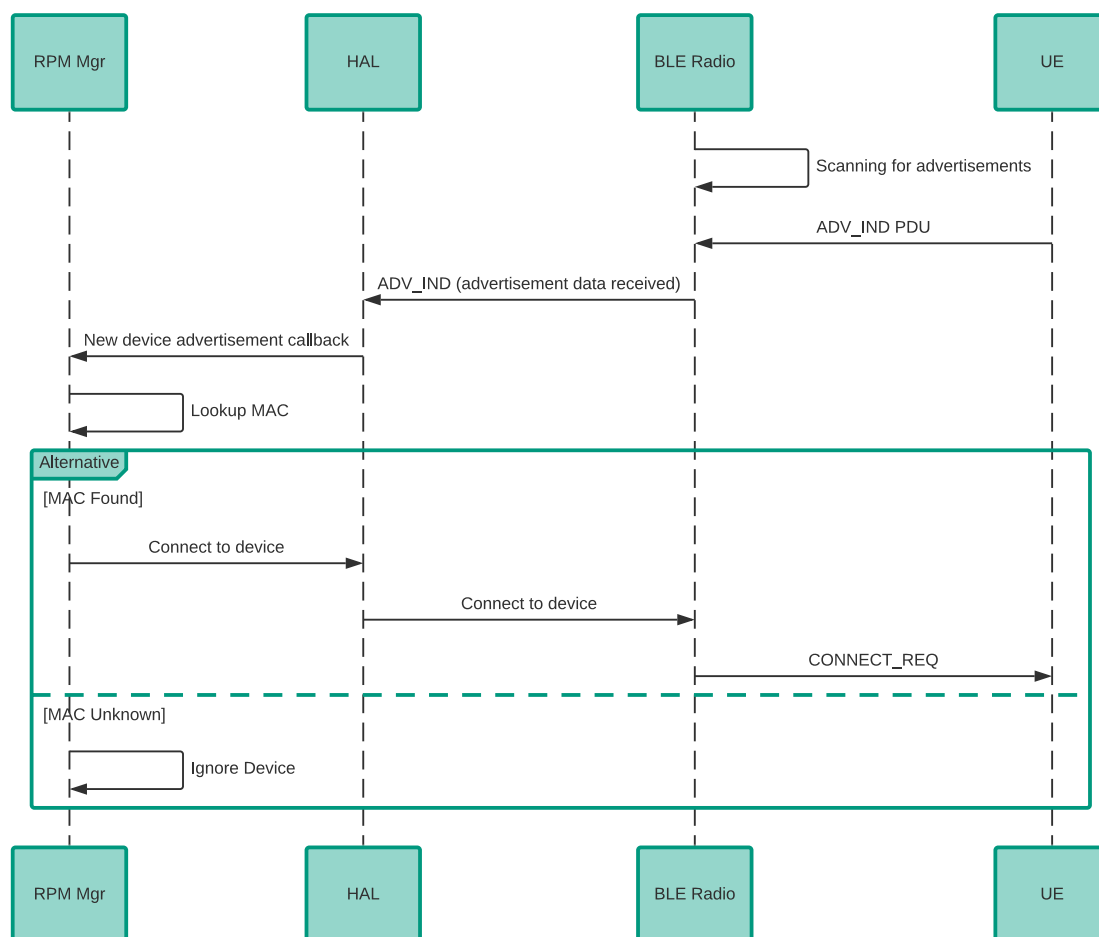
The system works by allowing applications to issue commands in response to events generated by the Bluetooth Low Energy hardware abstraction layer or by updates made to configuration tables in the synchronized database. The BLE events that were exposed for the proof of concept include: advertisement packet received, connection established, service discovered, characteristic discovered, notifications enabled, and new value received. The BLE commands that were exposed include: enable scanning, connect to device, discover services, discover characteristics, and enable notifications. This simple set of commands and events allowed for interactions with four different types of medical devices. They include a blood pressure cuff, pulse oximeter, scale, and thermometer. Each device was made by a different manufacturer.

Every device interaction begins with provisioning a specific medical device to a customer account through the RESTful API. The RESTful API exposes resources that are capable of performing database transactions. These database transactions are then replicated down to the Raspberry Pi via the magic of OpenSync. The information needed for provisioning the specific device is it's MAC address, the UUID of the GATT service it advertises, and the UUID of the GATT characteristic that stores the value of interest.

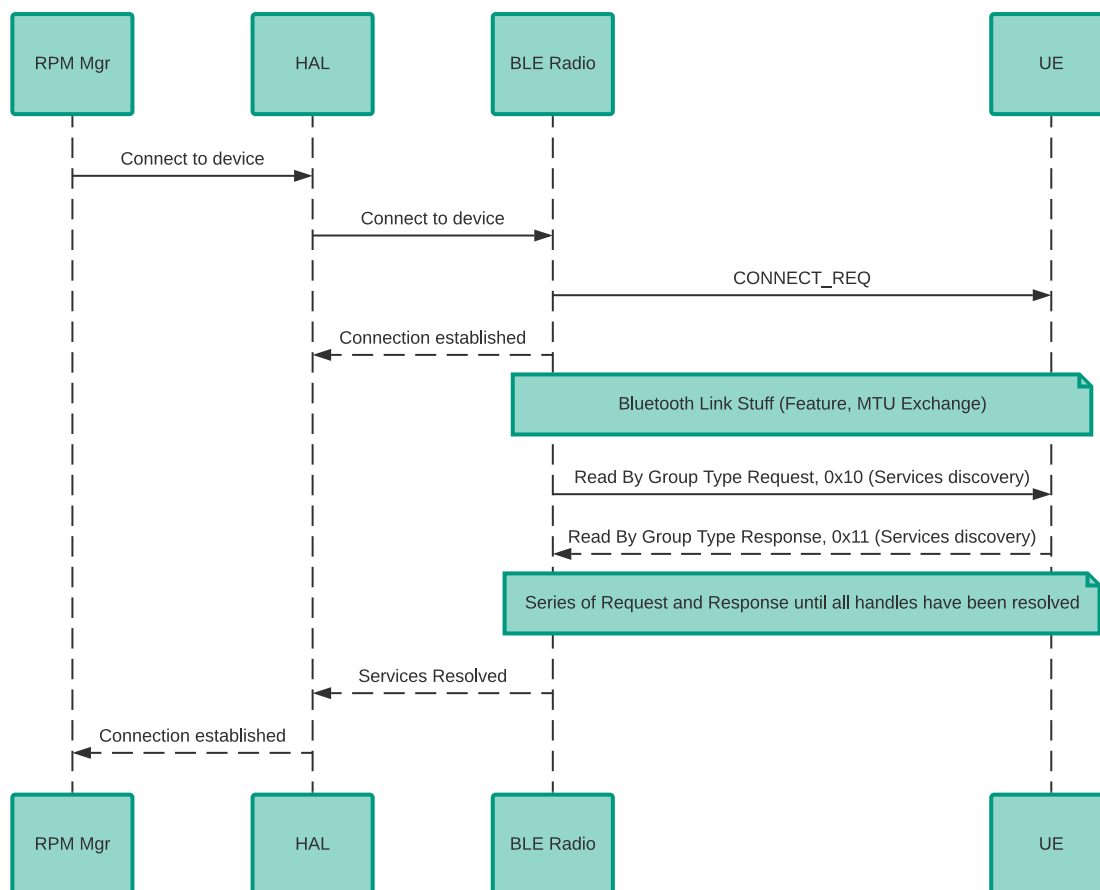
When the RPM manager code running on the Raspberry Pi detects the provisioning of a new device it enables scanning through the BLE HAL. The HAL then begins to pass received advertisement packets up to the RPM manager code. If the manager finds an advertisement packet with a matching MAC address it then issues a connect request to the device.

Upon successful connection the manager then begins a process of attribute discovery on the device. If one of the discovered attributes happens to match the UUID of the GATT characteristic value of interest, the manager issues an enable notifications command. The device will then begin transmitting measurements to the Raspberry Pi as they become available. When a new value is received the manager code takes the raw bytes and transmits them to the cloud via MQTT. Once in the cloud the data is then able to be ingested by any stakeholders that have authorized access.

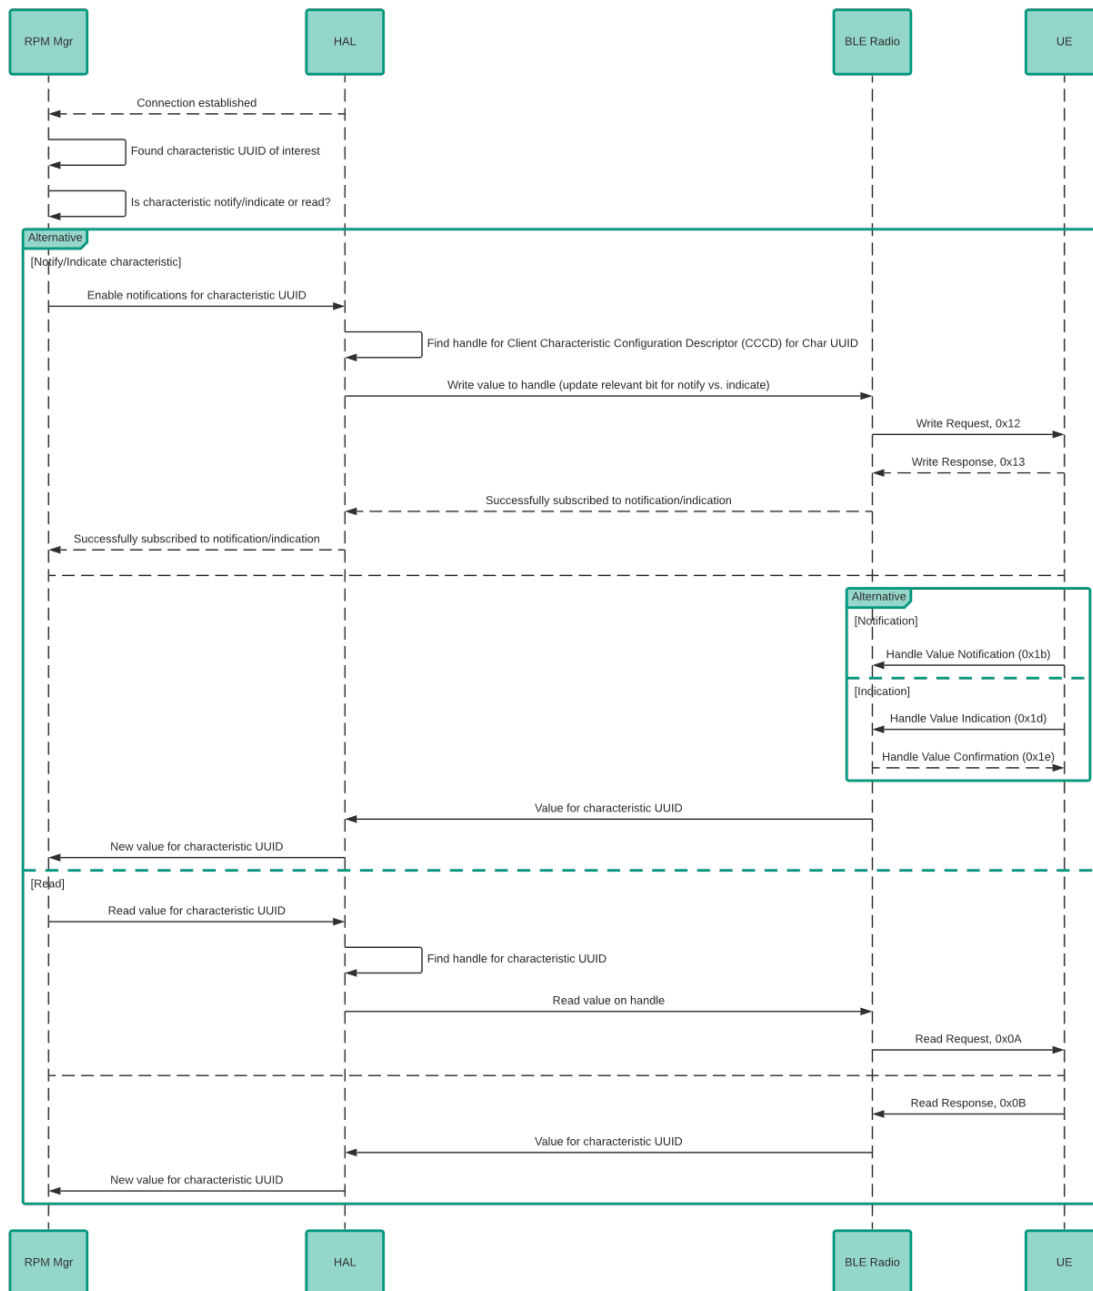
The main steps for getting data off the medical device can be broken down into three basic processes. The first process is establishing a connection between the medical device and router. The second process is discovering the attributes that the device exposes if this is not already known. The third process is enabling notifications and receiving values of interest. Below are sequence diagrams depicting the main processes that take place.



**Figure 1 – Device Discovery and Connection**



**Figure 2 – Device Connection and Attribute Discovery**



**Figure 3 – Data Transfer**

## **10. Areas Requiring Additional Work**

The RPM proof of concept demonstrated the bare minimum required to facilitate the transfer of data from connected medical devices. The above outlined approach does not address security. Utilizing the outlined components, security can be centrally administered, data can be guaranteed to be encrypted while in transit from the router to the cloud, and additional best practice security standards can be employed.

In addition to work done around security, a more robust set of interfaces to provision devices and ingest data should be designed. The simple user interface created for the proof of concept allowed for manual provisioning of devices to a single customer account and the ability to view measurements in real time. A more effective interface would allow the process of device provisioning and data access to be automated at the point of sale or when prescribed by a doctor. It should allow the customer to elect where their data should be transferred and should require consent for all third party access.

While the proof of concept shows that a router can be used for remote patient monitoring, it has not been tested with day to day use. Potential areas that can improve performance would be bonding with devices and maintaining device information such as mappings from a GATT UUID to an attribute handle. Additional systems can also be created to track battery levels, device connectivity history, and RSSI values.

## **11. Potential Concerns with the Outlined Approach**

The approach outlined above will work well for the majority of remote patient monitoring use cases. However, it suffers from two potential concerns. The first concerns placement of the router and signal range of the protocol. Bluetooth low energy has a range of 100 meters but in practical usage this is generally closer to 25 to 50 meters. In large homes or MDUs with lots of interference, the router may not be situated in an ideal location to accommodate these restrictions. In such situations placing IoT radios in Wi-Fi mesh access points would be an excellent approach. These can be small devices that create Personal Area Networks with a Wi-Fi data backhaul to the main router. In addition to providing extended IoT coverage they can also improve general Wi-Fi coverage.

The second concern relates to specific use cases requiring real time transmission of readings taken outside of a user's home. In this situation cable operators may be able to use a smartphone to serve as a hub for gathering and transmitting measurements with a companion app that complements the in home service. As Cable companies venture more into the mobile space this becomes a practical approach.

## **12. Conclusion**

Remote patient monitoring is an emergent technology in the delivery of healthcare that holds great potential. To fully capitalize on this potential a better platform for the transfer of data from connected medical devices is needed. The benefits of the approach described in this paper are more reliable connectivity, stricter security, cost efficacy, decoupling of data acquisition and data interaction, and an improved user experience. Cable companies are in a strong position to

leverage existing and enhanced infrastructure to support RPM technology. The rewards for implementing an RPM platform will be increased customer loyalty, new potential streams of revenue, and an opening to play a larger role in the broader smart home space. If Cable companies work together to create an open source standard many device manufacturers and RPM service providers will welcome the opportunity to use it.

## Abbreviations

|      |                             |
|------|-----------------------------|
| IoT  | Internet of Things          |
| BLE  | Bluetooth Low Energy        |
| PAN  | personal area network       |
| LAN  | local area network          |
| WAN  | wide area network           |
| RPM  | remote patient monitoring   |
| MAC  | media access control        |
| UUID | universal unique identifier |
| GATT | Generic Attribute Profile   |
| HAL  | hardware abstraction layer  |



# **Simultaneous Echo Cancellation and Upstream Signal Recovery using Deep Learning in Full-duplex DOCSIS Systems**

A Technical Paper prepared for SCTE•ISBE by

**Qi Zhou**

Georgia Institute of Technology  
75 5th St NW, Atlanta, GA 30308  
qi.zhou@gatech.edu

**You-Wei Chen**

Georgia Institute of Technology  
75 5th St NW, Atlanta, GA 30308  
yu-wei.chen@ece.gatech.edu

**Shuyi Shen**, Georgia Institute of Technology, Atlanta, GA

**Jeff Finkelstein**, Executive Director, Cox Communications, Atlanta, GA

**Drew Davis**, Executive Director, Cox Communications, Atlanta, GA

**Brian Lee**, Cox Communication, Atlanta, GA

**Gee-Kung Chang**, Professor, Georgia Institute of Technology, Atlanta, GA

# Table of Contents

| <b>Title</b>                                                         | <b>Page Number</b> |
|----------------------------------------------------------------------|--------------------|
| Table of Contents .....                                              | 2                  |
| Introduction .....                                                   | 3                  |
| Content .....                                                        | 3                  |
| 1. Echo Cancellation in FDX DOCSIS .....                             | 3                  |
| 2. Deep Learning Recap and DNN-based Echo Canceller Principles ..... | 5                  |
| 3. Experimental Setup and Results Analysis.....                      | 6                  |
| Conclusion .....                                                     | 9                  |
| Abbreviations.....                                                   | 9                  |
| Bibliography & References .....                                      | 10                 |

## List of Figures

| <b>Title</b>                                                                                                                         | <b>Page Number</b> |
|--------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| Figure 1 – Illustration of Self-Interference in FDX DOCSIS.....                                                                      | 4                  |
| Figure 2 – Conventional Echo Canceller Structure and Process.....                                                                    | 5                  |
| Figure 3 – Structure and parameters of proposed DNN echo canceller. ....                                                             | 6                  |
| Figure 4 – Comparison of DSP structure and time domain data management. (a) linear and nonlinear<br>cancellers (b) DNN decoder. .... | 6                  |
| Figure 5 – Experimental setup of the OFDM-based full-duplex transmission system.....                                                 | 7                  |
| Figure 6 – BER performance versus interference with (a) linear (b) nonlinear cancellers. ....                                        | 7                  |
| Figure 7 – BER performance verse interference with DNN decoder (a) train set (b) test set.....                                       | 8                  |
| Figure 8 – Wideband echo canceller BER performance (a) conventional DSPs (b) DNN. ....                                               | 8                  |

# Introduction

Full-duplex DOCSIS has been encountering a continuous uphill battle to double the channel capacity through resource sharing of uplink and downlink channels. As the downstream and upstream are delivered via the same spectrum at the same time, co-channel interference in the form of internal coupling, micro-reflections or echoes becomes a formidable challenge. To ensure the proper operations of full-duplex DOCSIS, echo cancellation is an urgently needed technique. Typically, it involves analog cancellation to lower the power level of the major echoes below the analog-to-digital converter (ADC) dynamic range, while digital cancellation is followed to remove the residual echoes ensuring a sufficient modulation error ratio (MER).

Many conventional echo cancellers are realized via a subtraction scheme, which implies the receiver is operating linearly. However, this limits the transceiver operation margin and the achievable MER. As the power of the desired and the echo signal increase, the receiver front-end may be driven away from its linear operation range and introduce nonlinear impairments. The crosstalk among the echoes and the desired upstream signals would degrade the echo cancellation performance. Thus, a simple subtraction-based cancellation is no longer sufficient.

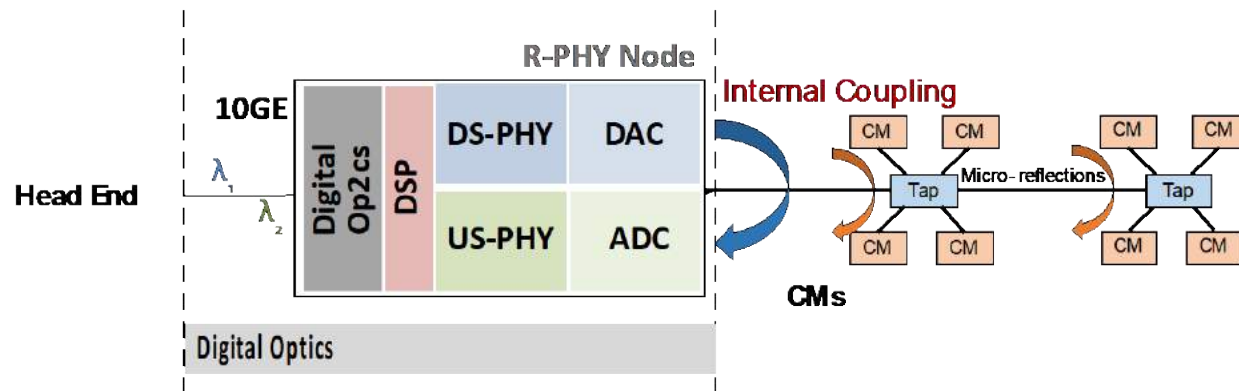
In this paper, we propose a deep neural network (DNN) based method to simultaneously cancel the echoes and recover the upstream signal. Both the received signal and the known downstream signal will be fed into the DNN processor, and the DNN output is the recovered upstream signal. The DNN is an efficient method to mitigate nonlinearities because of the implemented nonlinear activation function at each hidden layer. After proper initial training, the DNN-based canceller can achieve an excellent upstream signal recovery and outperform the conventional digital cancellation schemes. Moreover, the on-demand dynamic training can be performed implicitly without impacting the regular DOCSIS system operation. This novel approach would dramatically improve the recovered upstream signal quality and increase the capacity of the DOCSIS. This paper is organized as follows. Section 1 reviews the background of echo cancellation in full-duplex (FDX) DOCSIS and the conventional methods' limitation. Section 2 introduces the concept of deep learning and the principles of DNN to realize echo cancellation. Section 3 demonstrates a proof-of-concept experiment and results analysis of DNN based echo canceller.

## Content

### 1. Echo Cancellation in FDX DOCSIS

In the current cable access network, frequency division duplex (FDD) is implemented to provide various services to the cable subscribers. With FDD, the available cable spectrum is divided into non-overlapping parts for downstream and upstream, respectively. Different frequency splits have been defined including low split, mid split and high split. However, even with 5 MHz -204 MHz available spectrum of high split for upstream traffic from the cable modem (CM) to cable modem termination system (CMTS), it's still not sufficient to accommodate the exploding growth of bandwidth demanding services like AR/VR gaming, high-resolution video streaming. Besides, the high split also reduces the available spectrum to downstream traffic. To address the limited spectrum of cables and increase the spectrum usage efficiency, FDX was firstly introduced as DOCSIS 3.1 Full Duplex and latterly rebranded as part of DOCSIS 4.0 [1]. Unlike the frequency division multiplexing (FDM), the downstream spectrum could be reused by the upstream traffic without sacrificing the downstream bandwidth. In theory, the FDX DOCSIS allows overlapping spectrum between downstream and upstream traffic. Considering the 1.8 GHz full spectrum of the cable plant, the coax network bandwidth could potentially be doubled to 3.6 GHz. Though the appealing advantages of FDX DOCSIS, it suffers from various implementation challenges. One of the

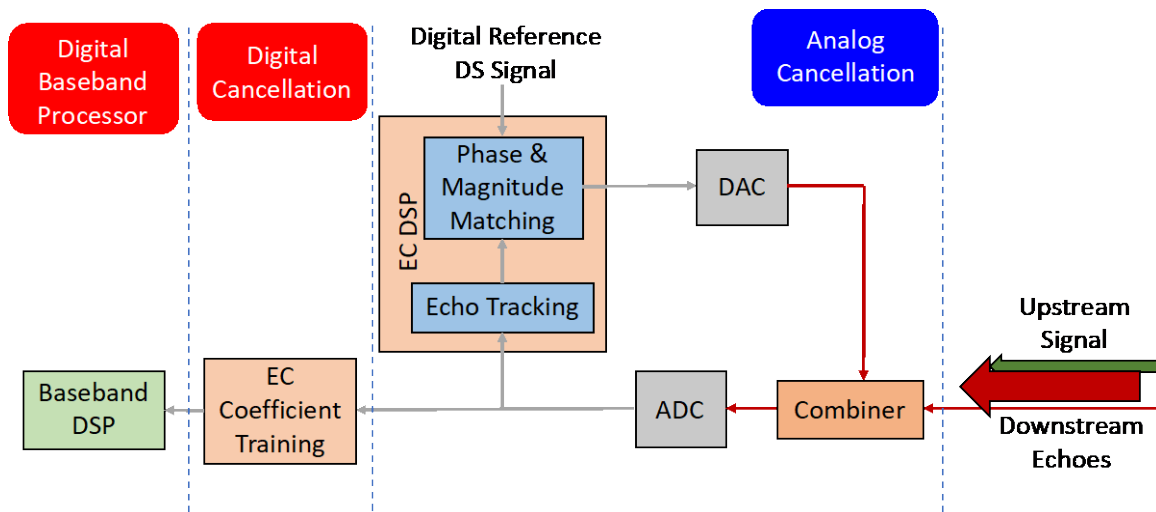
most significant challenge is the self-interference from the CMTS (Remote-PHY, RPD node) transmitter to its receiver. The system impairment due to self-interference is severe in FDX as the upstream and downstream share the same frequency spectrum, which cause co-channel interference. The self-interference arises from the internal coupling as well as the echoes in the forms of micro-reflections due to impedance mismatch at the taps as shown in Figure1.



**Figure 1 – Illustration of Self-Interference in FDX DOCSIS.**

Since the downstream signal has a much larger power compared with the upstream signal, the received upstream signal will be overwhelmed in the spectrum if there's no sufficient isolation. As the interfering down-stream (DS) signal is in-band, it cannot be directly removed by a radio frequency (RF) filter. Therefore, echo cancellation is necessary at the RPD-equipped node to cancel the self-interference and realize the FDX transmission in the cable network. It's typically assumed that N+0 network topology is necessary to support FDX DOCSIS as bidirectional FDX amplifier is associated with dramatic design implications [2]. N+0 topology imposes all passive RF components which can support bi-directional analog RF transmission based on Lorentz reciprocity theorem. Such that the FDX operation is possible without the expensive replacement on the already deployed cables and taps.

The conventional echo cancellation techniques in FDX DOCSIS split into two process, namely, Analog cancellation and digital cancellation. The logistics behind the split is to accommodate the limited dynamic range of the ADC at the node receiver. The power level of the reflections in the FDX spectrum can be more than 15 dB higher comparing to the desired upstream signal [2]. Applying analog cancellation before receiving by the node receiver can significantly relieve the saturation effect of ADC and improve the effective MER for the upstream signal detection. Figure 2 shows the typical setup and process for echo cancellation in FDX DOCSIS. The analog echo canceller takes a copy of the downstream signal and tunes its phase and magnitude to create a canceling signal with the same amplitude but with 180-degree phase difference. As the actual echo signals are from various sources and different locations, the multiple echo paths with different amplitude and phase are tracked and estimated in the digital domain and then converted into analog domain for cancellation using a DAC. After the analog cancellation is performed, digital cancellation is followed to further suppress the echoes using tracked/computed echo cancellation coefficients. The conventional echo canceller assumes that the self-interference (SI) can be removed by subtraction. However, this assumption doesn't hold true when there's nonlinearity at the receiver of the RPD. In realistic FDX implementation, an echo cancellation method which can mitigate the nonlinear impairments is necessary.




**Figure 2 – Conventional Echo Canceller Structure and Process.**

## 2. Deep Learning Recap and DNN-based Echo Canceller Principles

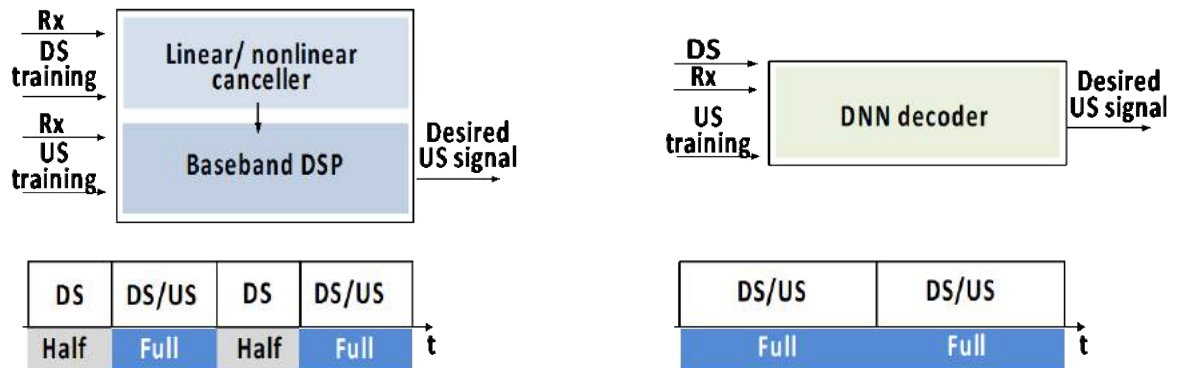
Deep learning has attracted tremendous interest from research and industry implementation. It is a sub-field of machine learning and it's based on artificial neural network with representation learning including supervised, semi-supervised and unsupervised. Deep learning performs excellently in image/video recognition, natural language processing and gaming. Beyond those conventional applications, deep learning also shows promising capability in dealing with challenges in wired/wireless communication, like multi-level signal recovery and adjacent channel interference mitigation [3-4]. Moreover, due to the booming development of neural network training algorithm, the training efficiency is significantly improved, while transfer learning can further reduce the required training time when there are any dynamic effects. Once the neural network is properly trained, the DNN inference is fast with low complexity as only simple matrix multiplications are required. Besides, the DNN can be implemented using general hardware like GPU, which provides the DNN with high scalability and low implementation cost.

To solve the aforementioned challenges in echo cancellation, we design and implement a DNN-based echo canceller to eliminate the nonlinear echo and the nonlinear crosstalk at the receiver jointly. The echo canceller will be implemented at the RPD. Here, we take the received signal and the known SI signal as the input to the DNN, while the output of the DNN is the recovered upstream signal of interest. The proposed DNN structure and parameters are shown in Figure 3, which consists of 3 hidden layers with 64, 32, and 16, respectively. The input layer has 41 neurons like the taps of a non-casual filter, while the output has 1 neuron representing the recovered signal-of-interests (SOI). The nonlinear rectified linear unit (ReLU) activation function is used at each hidden layer. In this case, the mapping between the received signal to SOI is not by subtraction of SI anymore. The nonlinear correlations between SI and SOI and the nonlinear distortion on the SI signals are also eliminated by the DNN after learning the sophisticated mapping. Besides, the DNN canceller realize the simultaneous echo cancellation and SOI recovery, greatly simplify the conventional two-step process.



**Figure 3 – Structure and parameters of proposed DNN echo canceller.**

Another benefit of the DNN canceller is its minimal impact on regular service during its training process. As shown in Figure 4, due to the limited modeling capability, the conventional digital signal processor (DSP) needs to be built in a block structure. The echo information is firstly estimated and removed from the received uplink signal with echo. After the echo canceller, a baseband DSP is conducted to recover the SOI subsequently. It is worth note that to obtain the accurate echo information, an additional training period is required in the conventional DSP fashion. In other words, certain time slots are preserved for DS signal transmission only, in such period, the upstream (US) signal is halted, and the DS training signal occupies all the available RE, resulting in a low channel utilization. On the other hand, the proposed DNN based decoder can break the DSP block structure. Most importantly, aided by the sophisticated DNN decoder, the dedicated training period can be removed which allows an implicit training without impacting system operation. Therefore, full-duplex transmission can be achieved all the time.




**Figure 4 – Comparison of DSP structure and time domain data management. (a) linear and nonlinear cancellers (b) DNN decoder.**

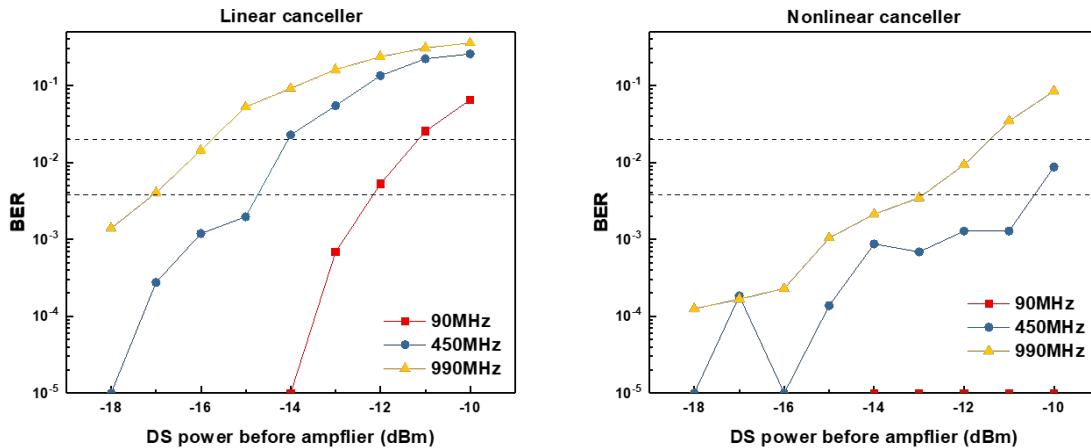
### 3. Experimental Setup and Results Analysis

The experimental setup of the orthogonal frequency division multiplex-based (OFDM) full-duplex transmission system is illustrated in Figure 5. A 16-GSa/s arbitrary waveform generator (AWG) is employed to generate the downlink echo and the desired uplink signal. To evaluate the impairment due to

the interference power from the reflected echo, a 10-dB attenuator cascaded by a 35-dB gain amplifier are applied to boost the downlink signal. Echo signal and the desired signal are combined via a power combiner before entering a 10 GSa/s real-time scope (RTS). Both echo and uplink signal are offline encoded and decoded via Matlab. The typical OFDM processing is employed. The FFT size is 2048 and the subcarrier spacing is set as 1.92 MHz. The bandwidth of echo is fixed as 990MHz and the desired uplink signal bandwidth under tested are 90, 450, and 990 MHz, respectively. The downstream transmitted power is ranging from -10 to -18 dBm before the power amplifier.



**Figure 5 – Experimental setup of the OFDM-based full-duplex transmission system.**



**Figure 6 – BER performance versus interference with (a) linear (b) nonlinear cancellers.**

As one can note that in Figure 6, the received BER performance is getting worse as the bandwidth of uplink increasing. This can be understood since the output  $V_{pp}$  of the desired signal, i.e., uplink, is fixed by the AWG as 1 V. Thus, as the bandwidth increasing, the per subcarrier SNR is decreased. Meanwhile, the BER is getting worse when the DS power increasing because the echo impairment is also increasing for the uplink signals in this case. On the other hand, the BER performance with nonlinear canceller, i.e., Volterra equalizer, always outperforms the linear canceller, i.e., minimum mean square error (MMSE), at the cost of a much higher DSP complexity. In the case of 90-MHz US, the BER of nonlinear canceller performs irrelevant to the power of echo interference.

The DNN decoding is a data driven processing. In this demonstration, the received signal is decoupled as training set, validation set and test set. Each set of data occupies one third of the received signal. The training set is firstly used to train the optimal weights and bias of the DNN, and the validation set is employed to validate the behavior of the DNN model. The DNN model we designed has one input layer,

three hidden layers, and one output layer. Activation function is  $ReLU(x)$  and the loss function is MSE. Optimizer employed in this experiment is Adamax which can ensure a faster and better convergence. After modeling, the test set would be processed by the trained DNN; while the training set result is also presented as a benchmark. It is worth to remind that DNN can cancel the SI and recover the SOI simultaneously, which enables full-duplex transmission.

As shown in Figure 7, the DNN decoding performance of 90 MHz US is similar to the nonlinear canceller. In this case, the DS information is not needed and thus all the RE can be fully utilized for the full-duplex operation. However, the performance of DNN drops quickly as the bandwidth increasing to 450MHz. DNN only outperforms the linear canceler in the strong interference range of DS power over -15 dBm, in such range, both DNN and linear canceler cannot achieve the FEC threshold requirement even with 20% overhead. This performance degradation may be caused by the insufficient vertical resolution of the received signal, since the peak to average power ratio (PAPR) increase proportional to the active subcarrier number. To circumvent this restriction, a straightforward method is reducing the active subcarrier number at the cost of reducing the bandwidth.

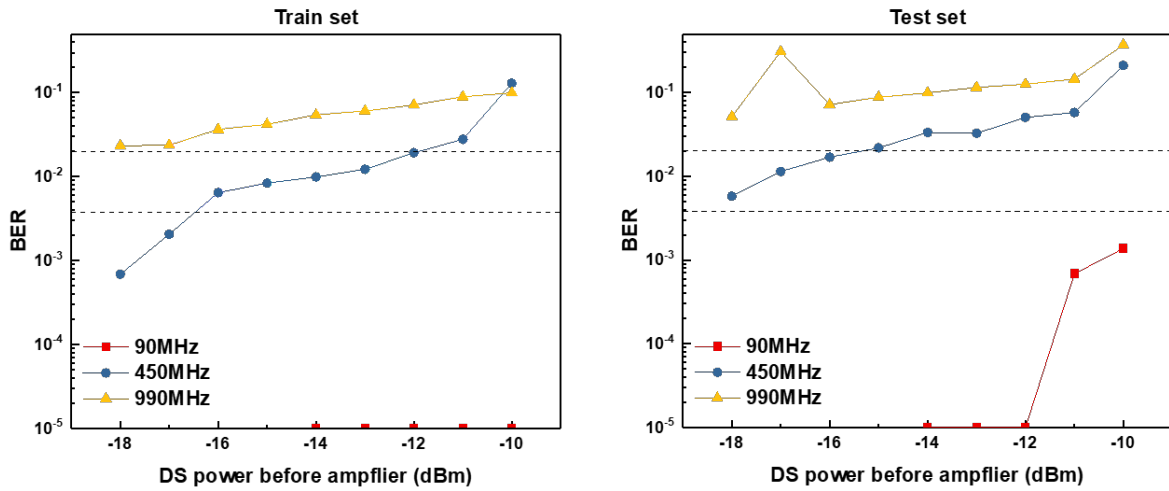


Figure 7 – BER performance verse interference with DNN decoder (a) train set (b) test set.




Figure 8 – Wideband echo canceller BER performance (a) conventional DSPs (b) DNN.



To get the better DNN performance and keep the transmission bandwidth, we can reduce the fast Fourier transform (FFT) size and active subcarrier number as well as increase the subcarrier spacing. Bandwidth under tested are 1GHz and 1.8GHz to comply with the DOCSIS 4.0 requirement. The active subcarrier number are 32 and 58, respectively. With the new settings, the US and echo have the same bandwidth for the FDX operation. the DNN canceller outperforms the conventional linear and non-linear DSPs in both scenarios as shown in Figure 8. The test set data performance is close to the training set, which implies no overfitting issue in the DNN process. The received performance is not sensitive to the bandwidth increment.

## Conclusion

As cable access network service providers evolve toward N+0 architecture and FDX DOCSIS for higher capacity, the echo cancellation technique becomes a core enabler to ensure reliable and high-performance services. The DNN-based echo canceler offers a future-proofing solution for cable operators to satisfy the exponentially growing bandwidth demand without dramatic infrastructure change.

In this paper, we explain the working principles of the proposed DNN canceller and conduct a proof-of-concept experiment to verify its performance in SOI recovery comparing with Volterra-based nonlinear canceler and conventional linear canceler. Better BER performance is observed at any measured DS transmission power for both 1 GHz and 1.8 GHz bandwidth cases. The results show the DNN canceller is robust for different configurations and working conditions of the coaxial cable network. This means that the cable operators may effectively support FDX operation over the whole available DOCSIS 4.0 spectrum without significant update on the physical infrastructure.

## Abbreviations

|        |                                                 |
|--------|-------------------------------------------------|
| ADC    | analog to digital converter                     |
| AWG    | arbitrary waveform generator                    |
| BER    | bit error rate                                  |
| CM     | cable modem                                     |
| CMTS   | cable modem termination system                  |
| DAC    | digital to analog converter                     |
| dB     | decibel                                         |
| dBm    | dB milliwatt                                    |
| DNN    | deep neural network                             |
| DOCSIS | Data Over Cable Service Interface Specification |
| DSP    | digital signal processing                       |
| DS     | down stream                                     |
| EC     | echo cancellation                               |
| EVM    | error vector magnitude                          |
| FDD    | frequency division duplex                       |
| FDX    | full duplex                                     |
| FEC    | forward error correction                        |
| FFT    | fast Fourier transform                          |
| GHz    | gigahertz                                       |
| HFC    | hybrid fiber-coax                               |

|        |                                               |
|--------|-----------------------------------------------|
| Hz     | hertz                                         |
| MER    | modulation error ratio                        |
| MHz    | megahertz                                     |
| MMSE   | minimum mean square error                     |
| MSE    | mean square error                             |
| PAPR   | peak to average power ratio                   |
| PHY    | physical layer                                |
| QAM    | quadrature amplitude modulation               |
| QPSK   | quadrature phase shift keying                 |
| R-PHY  | remote PHY                                    |
| RE     | resource element                              |
| ReLU   | rectified linear unit                         |
| RF     | radio frequency                               |
| RPD    | remote PHY device                             |
| RTS    | real-time scope                               |
| Rx     | received signal                               |
| SCTE   | Society of Cable Telecommunications Engineers |
| SD-FEC | soft decision forward error correction        |
| SI     | self-interference                             |
| SNR    | signal to noise ratio                         |
| SOI    | signal of interest                            |
| US     | upstream                                      |

## Bibliography & References

- [1] DOCSIS 4.0 Physical Layer Specification. <https://www.cablelabs.com/specifications/CM-SP-PHYv4.0>
- [2] H. Jin, and J. Chapman. "Echo cancellation techniques for supporting full duplex DOCSIS." In Proc. SCTE 2017, pp. 1-24. 2017.
- [3] Q. Zhou, F. Lu, M. Xu, P.C. Peng, S. Liu, S. Shen, R. Zhang, S. Yao, J. Finkelstein, and G.K. Chang. "Enhanced multi-level signal recovery in mobile fronthaul network using DNN decoder." IEEE Photonics Technology Letters 30, no. 17 (2018): 1511-1514.
- [4] S. Liu, Y. M. Alfadhli, S. Shen, M. Xu, H. Tian, and G.K. Chang. "A novel ANN equalizer to mitigate nonlinear interference in analog-RoF mobile fronthaul." IEEE Photonics Technology Letters 30, no. 19 (2018): 1675-1678.

# **Strategies for Implementing Edge Services in the 10G Cable Network**

A Technical Paper prepared for SCTE•ISBE by

**Eric Heaton**

Platform Solutions Architect

Intel Corporation – Network Platforms Group

2200 Mission College Blvd. Santa Clara, CA 95054

408-765-3447

[eric.d.heaton@intel.com](mailto:eric.d.heaton@intel.com)

# Table of Contents

| <b>Title</b>                                 | <b>Page Number</b> |
|----------------------------------------------|--------------------|
| 1. Introduction.....                         | 3                  |
| 2. Edge Deployment Models .....              | 5                  |
| 2.1. On-Premise Edge.....                    | 6                  |
| 2.2. Network Edge.....                       | 7                  |
| 3. Network Edge Platform Architectures ..... | 9                  |
| 3.1. CoSP + CSP Co-location.....             | 9                  |
| 3.2. CoSP Led + CoSP/CSP Services .....      | 11                 |
| 3.3. CSP Led.....                            | 12                 |
| 3.4. CoSP/CSP Aggregator .....               | 13                 |
| 4. Mapping to Real Estate.....               | 14                 |
| 5. A Converged Edge Architecture .....       | 16                 |
| 6. Considerations for Designing an Edge..... | 20                 |
| 6.1. A Summary of the Options.....           | 20                 |
| 6.2. Asking the Right Questions.....         | 22                 |
| Abbreviations.....                           | 23                 |
| Bibliography & References .....              | 24                 |

## List of Figures


| <b>Title</b>                                                                           | <b>Page Number</b> |
|----------------------------------------------------------------------------------------|--------------------|
| Figure 1 - Requirements on the Network for New Services.....                           | 3                  |
| Figure 2 - Use Cases and Associated KPIs.....                                          | 4                  |
| Figure 3 - Logical Parts of the Communication Service Provider Network.....            | 4                  |
| Figure 4 - Edge Locations and Terminology .....                                        | 6                  |
| Figure 5 - View of an On-Premise Edge Deployment .....                                 | 7                  |
| Figure 6 - View of a Network Edge Deployment .....                                     | 8                  |
| Figure 7 - Communications Service Provider and Cloud Service Provider Co-location..... | 10                 |
| Figure 8 - Communications Service Provider-Led Edge Deployment .....                   | 11                 |
| Figure 9 - Communications Service Provider-Led Edge Deployment.....                    | 12                 |
| Figure 10 - Communications Service Provider Aggregator-Led Edge Deployment.....        | 13                 |
| Figure 11 Type of locations in the MSO Network.....                                    | 14                 |
| Figure 12 - Network Edge Location Characteristics .....                                | 15                 |
| Figure 13 - Generic Access Platform with Compute Module .....                          | 16                 |
| Figure 14 - Access and Service Infrastructure in Silos.....                            | 17                 |
| Figure 15 - Converged Edge Architecture .....                                          | 19                 |
| Figure 16 - Mix and Match for an Edge Platform .....                                   | 20                 |

# 1. Introduction

Cable network bandwidth demands are growing exponentially as video becomes ubiquitous, Internet of Thing (IoT) devices proliferate, and new high bandwidth wired, and wireless Access technologies come online. Gartner estimates 90 percent of the data generated by the massive number of Internet-connected devices is sent to regional data centers for processing,<sup>1</sup> further stressing network infrastructure and increasing average response times for everyone.

That said, there is an incredible opportunity for broadband connectivity providers and those offering over-the-top (OTT) applications and Services to help make sense of and take action on the data coming from cars, cameras, factories, enterprises, and homes, and to do so in a timely manner. In fact, whole new categories of Services have been dreamed up, requiring ultra-low latency (i.e., augmented/virtual reality (AR/VR)), enhanced data privacy (i.e., medical records), or bandwidth optimization (i.e., video surveillance).

Figure 1 shows a range of industries and application segments that will benefit from these new network and compute capabilities if they can be delivered in a cost-effective manner.



**Figure 1 - Requirements on the Network for New Services**

For many networks, the latencies and other key performance indicators (KPIs) for specific Services within these segments, as shown in Figure 2, will require smart upgrades across the network infrastructure to:

- Reduce end-to-end latency by an order of magnitude
- Allow data to be processed closer to where it is generated, and
- Provide a coordinated deployment and management system to keep costs in line with revenue

<sup>1</sup> Gartner, "Edge computing promises near real-time insights and facilitates localized actions."; <https://www.gartner.com/smarterwithgartner/what-Edge-computing-means-for-infrastructure-and-operations-leaders>

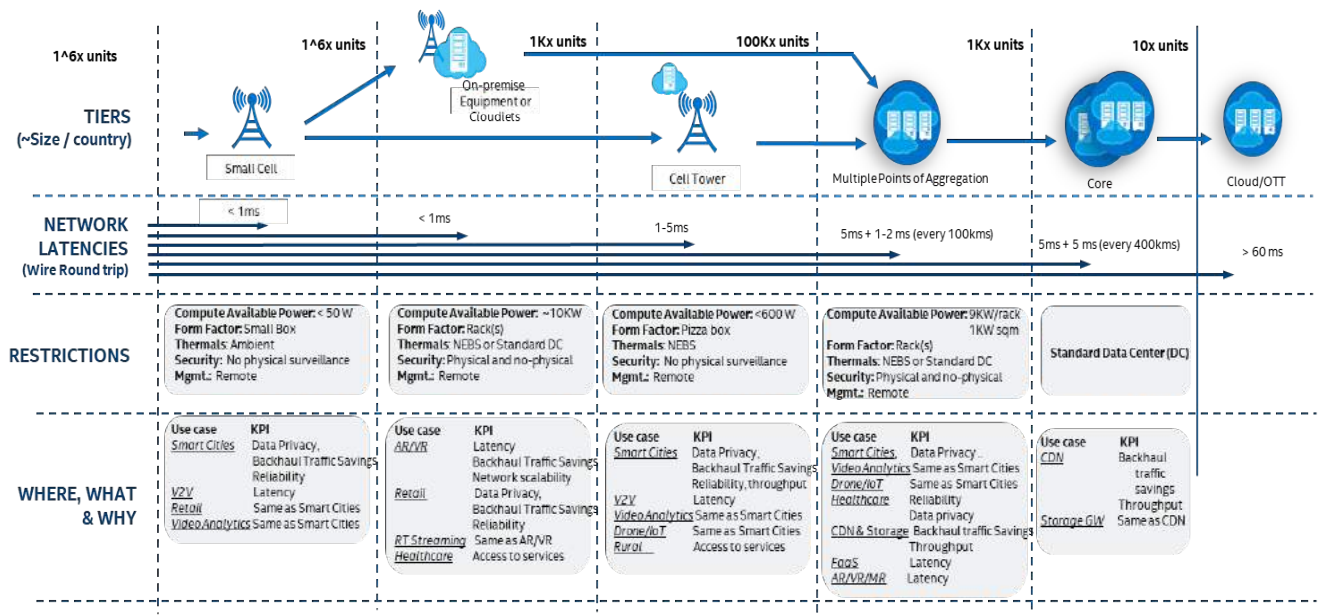


Figure 2 - Use Cases and Associated KPIs

Looking at the topology of the network, "The Edge" makes for an interesting and obvious place to manifest key infrastructure because it is physically closer to end users. The reality is there are multiple places that could be considered the Edge, and they all can exist in the same network. Figure 3 shows that at the top level, the Edge can be split into an On-Premises Edge and a Network Edge. The next section will go deeper into defining both deployment models, what types of Services are best provided by one model compared to the other, and other considerations to deliver maximum flexibility and return on investment (ROI).

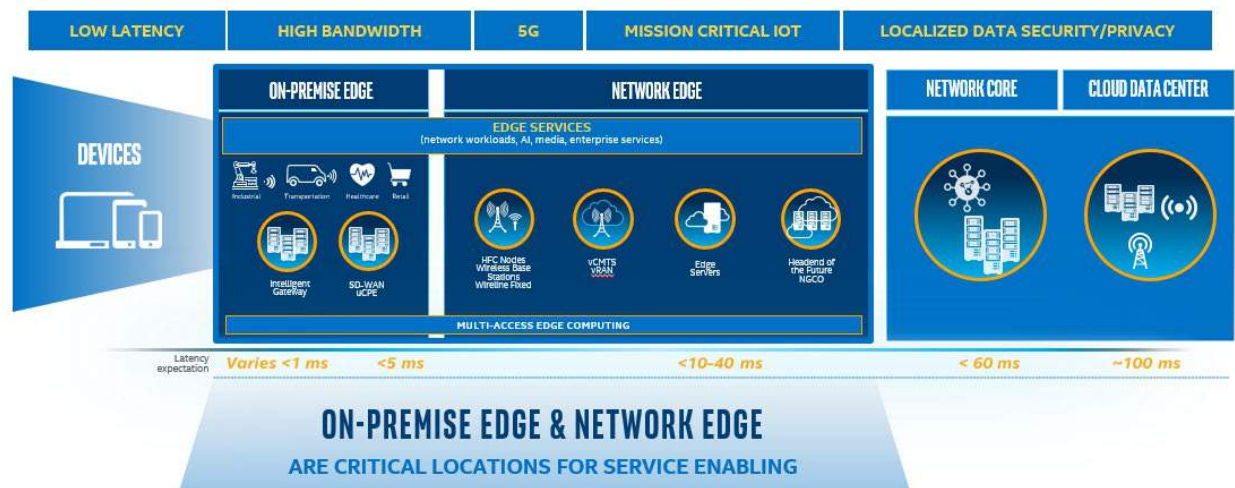


Figure 3 - Logical Parts of the Communication Service Provider Network

Figuring out where network infrastructure may be hosted is only part of the battle. This paper will continue on to discuss considerations for designing and buying Edge platforms, as in the actual hardware and software that will run the network, how functions are split across different equipment, and who owns the functions, users, and traffic.

In a single-Access or single-Service world, it is easy to line up bespoke solutions that include appliances and custom management interfaces. Perhaps the finances work out for two or three solutions set up in parallel. But to scale Edge infrastructure and maximize resources and operational efforts, a common platform (or at least common building blocks) that addresses all Access technologies and Service deployments will make the most sense. To this end, this paper introduces the idea of a Converged Edge Architecture, comprised of commercial off-the-shelf (COTS) equipment and standard software frameworks and interfaces that can be used to develop any number of specific solutions. This approach reduces complexity and time-to-market and allows for common orchestration across any Edge location. Converged Edge Architecture does not define a single platform, but rather is a common framework that will deliver the right Edge platform for the functional needs and environmental constraints of a given location in the network.

There are many reasons the cloud is moving to the Edge, including network optimization through virtualization and data locality, cost savings through white box platforms and automation, and new monetization opportunities through the introduction of new Services and business relationships. This paper primarily focuses on the last category – opportunities and considerations around implementing Services – though topics like virtualization and automation will be referenced as they underlie modern network architecture and implementation.

In fact, recent headlines<sup>2</sup> have shown that some communications service providers (CoSPs) and cloud service providers (CSPs) are already implementing Services at the Edge and selling them commercially. This paper should serve as further encouragement and also as a guide to show there are many considerations to creating an effective and scalable Edge that supports both multiple Access technologies and the latest Services, hosted by the CoSPs, CSPs, and companies with OTT offerings. It is the right time for multiple system operators (MSOs) to plan their new Edge(s), as it can be done in conjunction with the ongoing march towards a Distributed Access Architecture and a virtualization environment in its many forms.

By asking the right questions and knowing some of the key architectural options discussed in this paper, network operators – and MSOs in particular – will be in a better position to realize the value of their existing infrastructure and improve customer experience in an increasingly Edge-focused world.

## 2. Edge Deployment Models

Figure 4 breaks down the taxonomy of Edge platforms and locations beyond the top-level of an On-Premises or a Network Edge. The On-Premises Edge, or “On-Prem Edge,” can be broken down into platforms/locations for smart sensors, intelligent gateways, and intelligent Edge servers, with generally increasing complexity, power, and capability as one moves left to right. Similarly, the Network Edge contains platforms and locations at the Access Edge, near Edge, and data center Edge.

Note that some networks may not manifest all these categories because the physical real estate to host certain equipment might not exist, locations and functions may have been consolidated onto fewer

---

<sup>2</sup> Fierce Wireless, “Verizon, AWS bring 5G MEC to Boston, Bay Area”; <https://www.fiercewireless.com/operators/verizon-aws-bring-5g-mec-to-boston-bay-area>

platforms, or some functions may not be implemented for one reason or another. Conversely, the network operator does not have to necessarily choose On-Prem Edge types or Network Edge types exclusively, as they each have their own pros and cons with respect to achieving different types of technical and business goals.



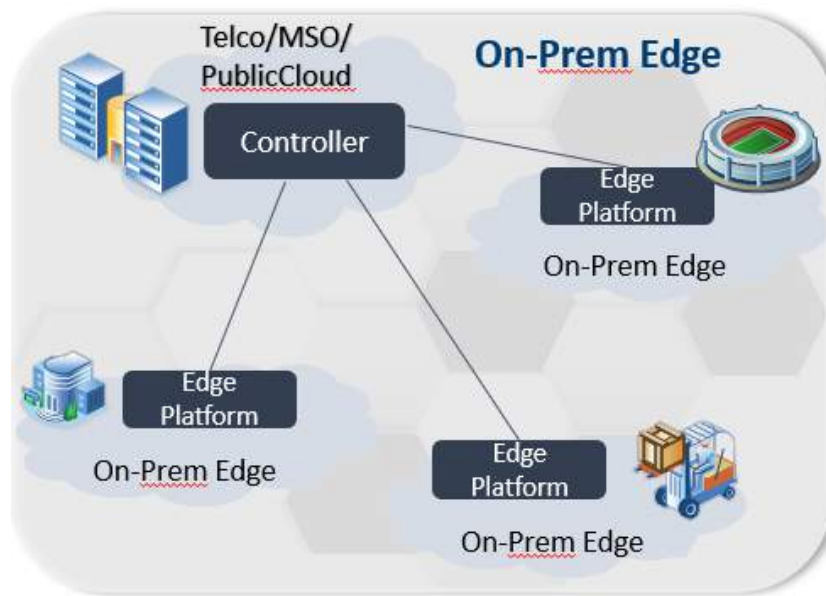
**Figure 4 - Edge Locations and Terminology**

The following sections cover these deployment models in turn.

## 2.1. On-Premise Edge

Figure 5 shows an On-Prem Edge-type deployment model. Here, there is a controller for Services that is located at some centralized place in the network and manages functions and Services that ultimately run on an Edge platform, like a universal customer-premises equipment (uCPE). Non-real-time functions such as controllers and Service management will usually reside in the most cost-effective place, which is typically deep in the network at a regional data center or even in a public cloud. That said, the controller may absolutely be deployed at another Network Edge location, like a Headend or Hub, to comply with legal requirements, or perhaps to satisfy operator or customer requirements for full data locality.





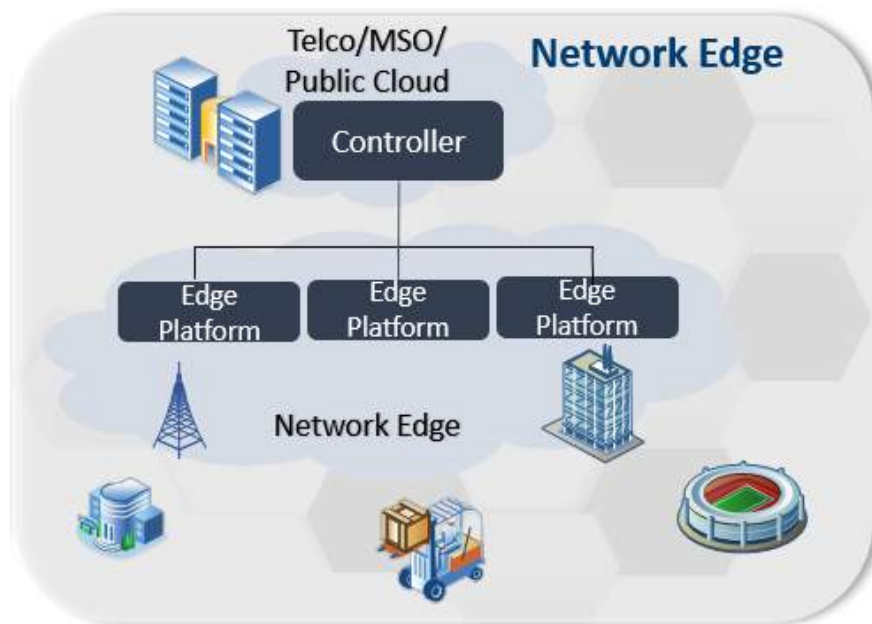
**Figure 5 - View of an On-Premise Edge Deployment**

Growing at a triple-digit rate in recent years,<sup>3</sup> SD-WAN solutions are generally offered through an On-Prem model, serving various applications for stadiums, farms, industrial IoT, and the like. True to its name, the main distinguishing factor for the On-Prem model is it includes a flexible platform at the customer premises that can run dynamic, real-time workloads locally. These workloads can be virtual network functions (VNFs) or Services.

While this paper focuses more about the requirements and options in the Network Edge, the platform choices for On-Prem are comprehended in the Converged Edge Architecture framework introduced in Section 5.

## 2.2. Network Edge

Figure 6 shows a Network Edge-type deployment model. In this case, there is also a controller – or more likely, a set of controllers – within the CoSP hierarchy to control network function virtualization infrastructure, multiple VNFs, and/or Services hosted directly by the operator or its partners. These controllers will orchestrate and manage such functions to run on Edge platforms located generally at Nodes, Hubs, Headends, point-of-presence server locations, or central offices. As in the On-Prem model, the controllers could exist anywhere “upstream” from Edge platforms in the network; it is likely that they will run from a regional data center, private cloud, or at a CSP.



**Figure 6 - View of a Network Edge Deployment**

This infrastructure will support broadband Access, perhaps more than one type, as well as host Services that can be delivered by operators themselves, CSPs, or third parties. Given all these possibilities, the Edge platforms in Figure 5 are not necessarily a single server or appliance but can manifest as a collection of equipment at a location, providing a set of APIs upstream to controllers and other functions, and downstream to users. In short, Edge platforms may involve more than one piece of hardware or more than one physical location.

To make the most of a Network Edge deployment and to keep costs down as the number of Access types and Services increases, it will be key to find as much common ground as possible between the different hardware elements and to provide a homogeneous software layer for management. These are some of the main goals of the Converged Edge Architecture effort covered later.

Many of the platform-as-a-service products that have been announced<sup>3</sup> by various CoSPs, CSPs, Edge compute specialists, and real estate management companies (i.e., towers) are based on a Network Edge deployment model. Behind the scenes, there are a variety of ways these parties can organize themselves and their resources to make their own Edge products. They will ask themselves – what network functions and Services should we offer (and who should own these products), where should the equipment exist, who should own the customers, how should revenue be divided, and so forth. The answers to these questions have both business and technical repercussions, and for the latter, will affect how the Edge platforms/locations can be architected. The next section describes several Network Edge platform architectures emerging in the marketplace and the associated advantages and disadvantages from the network operator’s perspective.

<sup>3</sup> RTTNews.com, “Microsoft To Use Telefonica Infrastructure for Datacenter Region in Spain”; <https://www.nasdaq.com/articles/microsoft-to-use-telefonica-infrastructure-for-datacenter-region-in-spain-2020-02-26>

### 3. Network Edge Platform Architectures

The On-Prem network deployment model has its own implementation challenges, but at a high level, network functions and Services will run on a uCPE at the customer premise with a software controller running deeper in the network on COTS servers and managing those Services.

As discussed earlier, the Network Edge case has many options for splitting up functions across different network locations and deciding who will own those functions and equipment (i.e., some may be owned by a CoSP, CSP, or a third-party). The CoSP Edge architect will have to consider which Services the network operator would like to host, where equipment (and its capabilities) can be deployed, how to split network functions (i.e., controllers versus data plane), and who is going to own which parts of the solution in order to come up with a comprehensive Edge platform architecture.

There are several divergent approaches emerging in the market. Each approach allows for different types of business arrangements (i.e., how is revenue paid and split up) and will require different technical arrangements to be made between the partners involved.

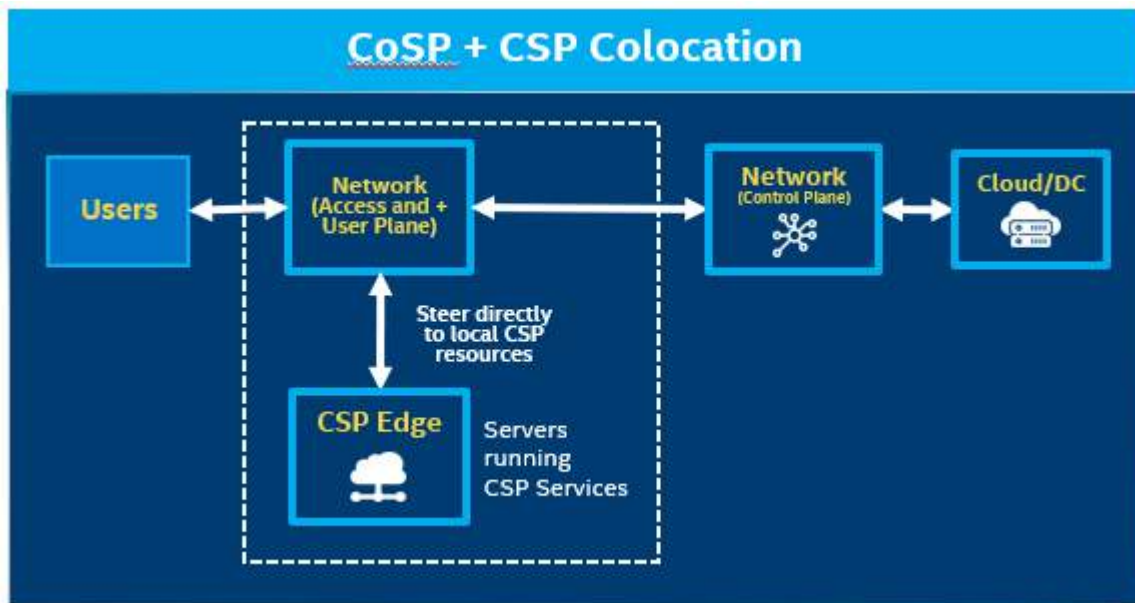
1. CoSP + CSP Colocation
2. CoSP Led
3. CSP Led
4. CoSP Aggregator

The following sub-sections will review what it takes to implement each one of these business arrangements, along with the main driving forces for a CoSP to pursue one over the over.

#### 3.1. CoSP + CSP Co-location

Currently, the most popular, or at least the most talked about, architecture for a Network Edge deployment is for a CSP to co-locate equipment to deliver their Services at Edge locations owned by the CoSP. Figure 7 shows the logical view of how this would work. Presumably, the locations being “shared” by the CoSP and CSP would have latency advantages over what the CSP could promise on their own, with the CoSP having physical real estate very close to end users.

A CSP could then charge a premium when offering a content delivery network (CDN), a video analytic engine, or a generic platform-as-a-service with tighter and better guarantees around latency or data locality than the same Services being deployed at a regional data center or at their own site.



**Figure 7 - Communications Service Provider and Cloud Service Provider Co-location**

The most straightforward benefit for the CoSP is it gets some sort of “rent” or revenue share from the CSP. More interesting for business development is that these partners can engage in joint marketing to advertise the CSP is offering Services over the Access medium owned by the CoSP that would not be available through other means – a “better together” story. There may also be practical reasons to enter into this type of agreement as the CoSP may not have the expertise or desire to develop and manage whatever Services are being provided by the CSP.

From the CoSP perspective, the downside is that it does not have direct control on how co-located Services are monetized, and, further, would not have access to the user information or traffic telemetry for secondary opportunities for data monetization. In addition, there may be some inadvertent CoSP/CSP lock-, or at least constrain the relationship to one CSP at a time if the CoSP location is not set up to easily host and secure multiple parties.

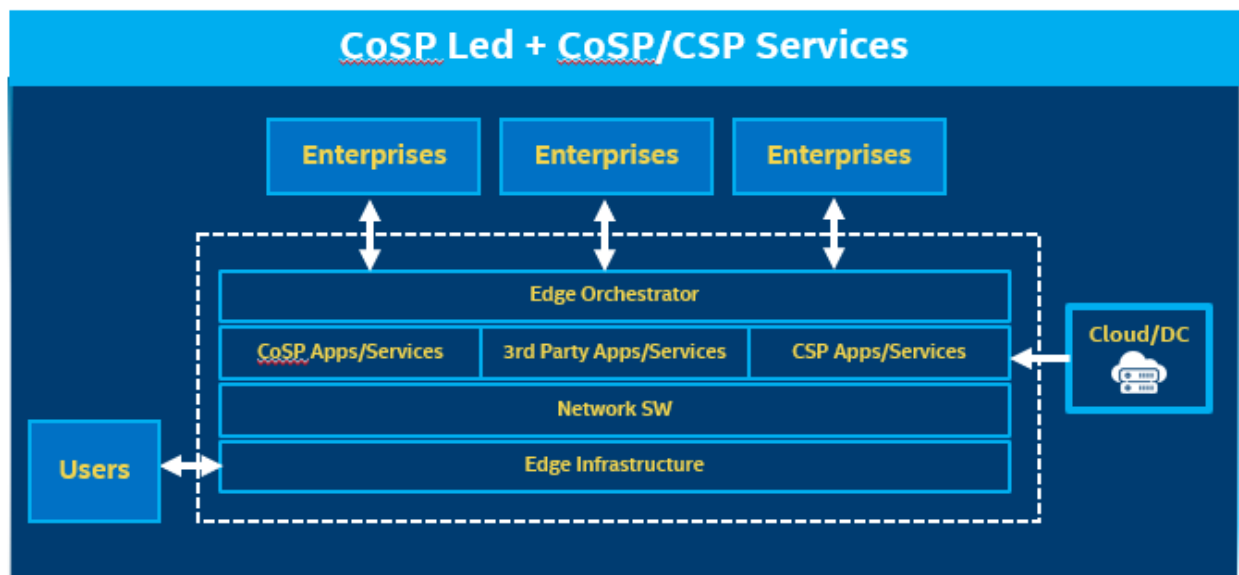
For this model, the “Edge platform” is made up of two subsets of equipment hosted at the same Headend, Hub, or Node: one owned by the CoSP providing the network connectivity and the other owned by the CSP providing the Services. The CoSP and the CSP would select their individual equipment based on their own experiences and needs.

Note that the CoSP could also become a customer of the CSP and have them host new Services on their behalf, so this arrangement does not preclude the CoSP trying their hand at offering Services. It’s just that the Cloud/Service infrastructure is being provided by CSP and getting paid for it. In this case, though, perhaps the co-location aspect of this arrangement would even allow knowledge transfer of Cloud technologies to the CoSP?

### 3.2. CoSP Led + CoSP/CSP Services

The second most common architecture is for the CoSP to own the Edge platform altogether, hosting CSP Services on it on behalf of the CSP. In this case, the CoSP must develop an API, allowing CSPs to access resources at advantageous Edge locations. With this architecture, the CoSP owns the execution and delivery of the Services to the customer. Figure 8 shows the logical diagram of the platform.

From the CoSP perspective, there is more complexity and internal expertise needed, as they need to develop and maintain their own commercial platform-as-a-service. CoSPs do not have ownership of the Services themselves but would be able to monetize the management of the Edge locations and equipment.



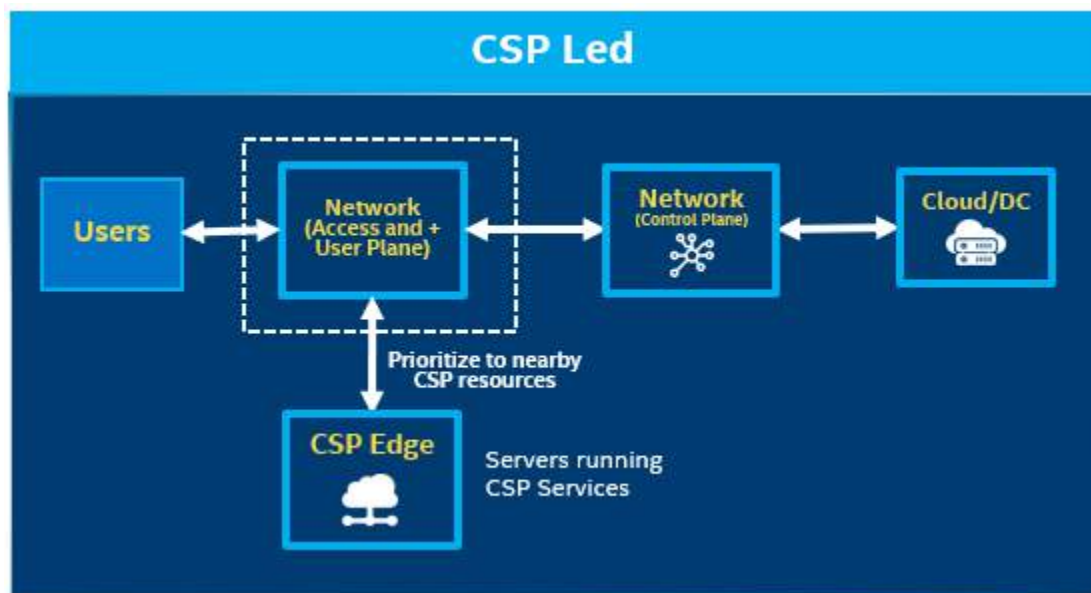
**Figure 8 - Communications Service Provider-Led Edge Deployment**

The additional benefits for the CoSP is it does not have to deal with the logistics of hosting another company's equipment in their buildings, and it has a unified interface to sell its Edge platform to any number of CSPs or third parties. Developing this model also implies the CoSP invested in internal resources that understand cloud technologies, which is generally a good thing as more elements of the network are virtualized.

For this model, the Edge platform is wholly specified and managed by the CoSP. There can still be multiple types of equipment involved in the final implementation, but a single owner is more likely to host all the different functions and Services on a common set of hardware in order to get better economies of scale for both the capital costs and for the ongoing costs towards the workforce. High performance VNFs running on COTS servers and new options for programmable switches are making it easier to reduce the number of specialized appliances in the network, as described in more detail in the next section.

### 3.3. CSP Led

The last two models to be presented may be less prevalent, but nonetheless fit certain niches. In the CSP-led model shown in Figure 9, the CSP owns the Edge locations and platforms. This assumes the CSP has real estate close enough to users to distinguish their Service offerings from those coming from the general cloud. In this case, the CSP uses a CoSP for last mile connectivity, but otherwise owns all aspects of the Service delivery and management.



**Figure 9 - Communications Service Provider-Led Edge Deployment**

For the CoSP, this model may just be construed as a high-end business as usual arrangement, where a CSP or any other customer is paying for a broadband offering, albeit one with low or ultra-low latency guarantees. That said, CoSPs are finding that partnerships with CSPs focused on Edge-specific infrastructure (i.e. that have developed, or can develop, many points of presence near last-mile CoSP locations) may actually make for a convenient alternative to the CoSP + CSP Co-location option insofar that a similar “better together” story can be created but without the added logistics and coordination of sharing the same physical space and Headend or Hub management<sup>4</sup>.

In this model, the CoSP and CSP platforms can be developed and deployed relatively independently if the connection between the two networks can minimize latency to an absolute minimum.

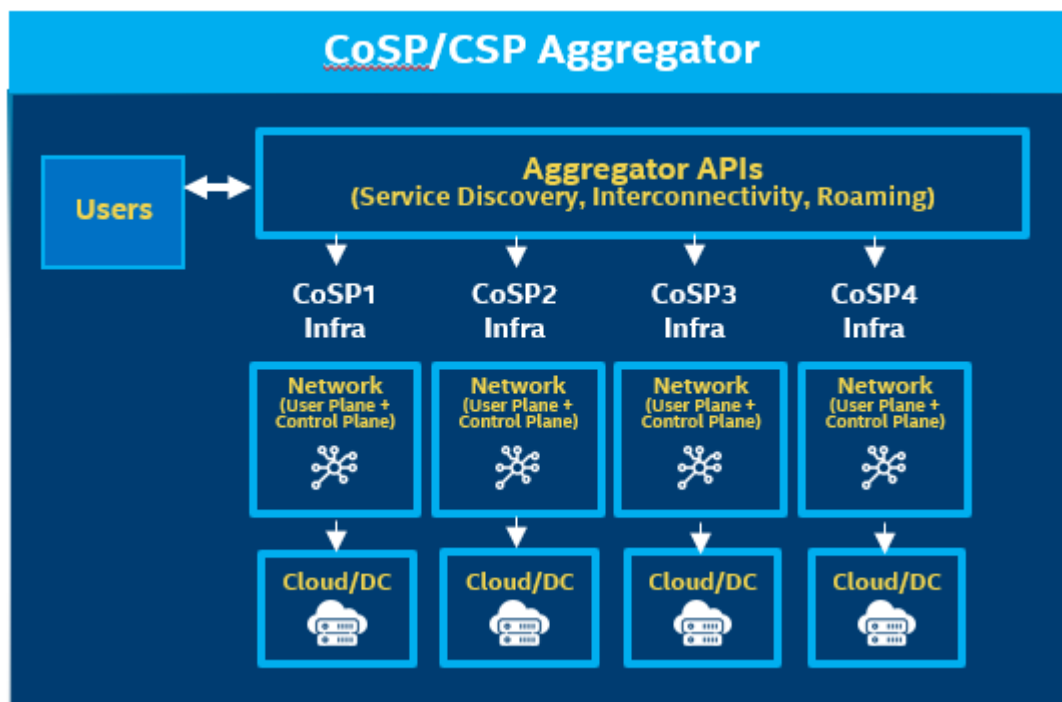
So, the CoSP’s Network Edge platform is whatever the CoSP wants to use for its most advanced Access and broadband offerings and should follow the best practices of the industry (i.e. generally moving to DAA and virtualization-based solutions). Similarly, the Network Edge platform of the CSP will be optimized for the products and Services they are offering and generally will be a common software management infrastructure hosted on COTS servers and switches – typical for any CSP.

<sup>4</sup> Robuck, “Cox targets the Edge for the next evolution of network performance and security”, <https://www.fiercetelecom.com/telecom/cox-targets-Edge-for-next-evolution-network-performance-and-security>

Like the CoSP/CSP Co-location model, the CoSP could also become a customer of the CSP, hosting new Services on their behalf.

### 3.4. CoSP/CSP Aggregator

This last model is one in which a third-party aggregates connectivity and Service options from a variety of CoSPs and/or CSPs and offers a common API to other service providers. In this case, the aggregator owns its own real estate, buys connectivity from one or more CoSPs, contracts with CSPs for Services, and perhaps even offers its own Edge Services. The result is an aggregated Edge offering.<sup>5</sup> Figure 10 shows how such an aggregator can develop its own platform or Service API that then plugs into the offerings of partner companies for execution and delivery.



**Figure 10 - Communications Service Provider Aggregator-Led Edge Deployment**

CoSPs could monetize this arrangement in multiple ways. If they are just selling broadband connectivity, then this basically is the same as the CSP model. Alternatively, they could move up the food chain and offer their own platform-as-a-service, as described in the CoSP-led approach mentioned previously.

For the aggregator, the Access and broadband connectivity from CoSPs, and therefore the Services coming from CSPs, is ideally via an Ethernet/IP network. In this case, the Edge platform architects do not have to worry about the disparate requirements (and sometimes baggage) of physical Access technologies and therefore it is easier to choose COTS hardware – servers and switches – for the reasons mentioned in previous sections.

<sup>5</sup> Dano, “SBA, American Tower double down on Edge computing opportunity”; <https://www.lightreading.com/the-Edge/sba-american-tower-double-down-on-Edge-computing-opportunity/d/d-id/762941>

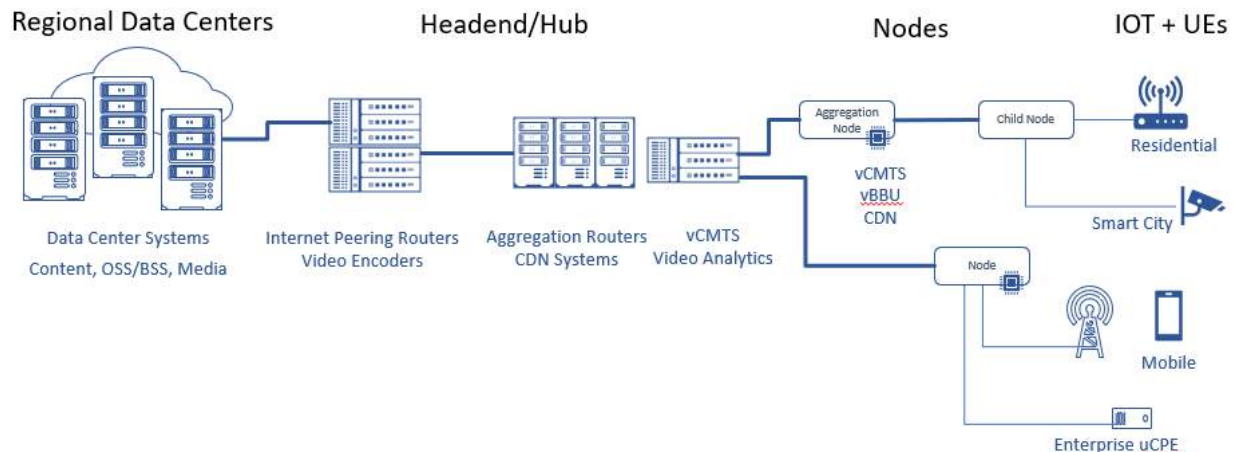


## 4. Mapping to Real Estate

This paper has discussed both Edge deployment models and Edge platform architectures as logical entities, but in real life, the equipment, software, and operations to manifest these things need to live in physical locations in the MSO network. This section maps the theoretical to the empirical world and discusses the competitive advantages cable networks may have compared to Access technologies.

CableLabs recently described typical MSO network locations and how their characteristics could apply to Edge deployments.<sup>6</sup> Figure 11 shows how the following locations are connected and how far they are from end users:

- Cloud
- Regional Data Centers
- Headend/Hub
- Aggregation Nodes
- Child Nodes
- Customer Premises
- Cell Sites



**Figure 11 - Type of locations in the MSO Network**

On-prem or Network Edge locations have associated benefits and constraints. On the benefits side, the closer one gets to the IoT devices and users, the lower the latency, the higher the data locality, and the lower bandwidth required upstream in the network. On the constraints side, the closer one gets to the IoT devices and users, the less power, the less environmental control, and the more costly it is to deploy and service equipment.

The benefits for On-Prem or Network Edge are straightforward to understand – “closer is better” for latency and data locality. The following discusses the constraints for the different types of host equipment that can be hosted.

<sup>6</sup> Levensalor, Stuart, “The Modular, Virtualized Edge for the Cable Access Network”, <https://community.cablelabs.com/wiki/plugins/servlet/cablelabs/alfresco/download?id=2c46cef2-af44-47be-bdd4-98a948cbc60d>



The left side of Figure 12 shows a large, multi-story, regional data center (e.g., central office or CO) or large Headend, serving tens or hundreds of thousands of users, and for the most part, the set up can use typical data center approaches. Moving to the right along this continuum towards the users, the environmental constraints increase, the compute capacity goes down, and the deployment and management of equipment is more costly.



**Figure 12 - Network Edge Location Characteristics**


CoSPs need to decide if they want all the above locations to be “available” for their Edge. That is, will a given location and the equipment therein not only have the right level of physical connectivity in both directions, but will they also have an ability to run network functions or enterprise software? Further, how much of this hardware and software can be brought into a centralized management domain?

In the ideal cloud-extended-to-the-Edge vision, every location is part of a large pool of flexible, distributed, compute, storage, and networking resources and functions; Services can be set to run where they are needed to satisfy technical and business needs at the lowest possible cost. This is easiest when all the hardware is common, and the software has similar resource needs – as in data center and CSP locations. However, the further out one gets from the regional data centers, the more likely software workloads have higher data plane needs, lower latency requirements, and traditionally, have been served by specialized appliances. These types of devices will be part of an Edge plan for many operators, noting that fixed-function devices raise the TCO, and by definition, limit the flexibility of the solution.

As discussed earlier, the move to Distributed Access Architecture (DAA), virtualization, the power of general-purpose compute via COTS servers, and the growing market for programmable switches will reduce the need for such legacy solutions. Certainly, many Edge platforms residing in regional data centers, Headends, and Hubs will be based on COTS servers and switches.

Regarding the locations in the outside plant, it may not have been possible from a technical or economic standpoint to have flexible compute resources at the Nodes or smaller huts and cabinets in the past, but this is changing. Operators still need to do their homework to understand what is possible based on their

existing infrastructure and what types of businesses they want to pursue in the future, but the expanded availability of NEBS-compliant servers and industry innovations, like the SCTE Generic Access Platform (GAP), will bring down the cost of putting small-form-factor servers at these types of locations.



**Figure 13 - Generic Access Platform with Compute Module**

Figure 13 shows an example of a GAP-compliant Node in which the form factor, the electrical and logical connectivity, and module management are standardized, such that a vendor can provide the same types of compute, storage, and network capabilities found in a data center. The scale might be different – in the hybrid fiber-coaxial (HFC) case, the Node may serve only a couple of service groups – but the architecture and the way its resources can appear to the larger management infrastructure of the network operator are very much the same.

What developments like GAP mean is that it is now possible, from both technical and business perspectives, to distribute flexible resources to all parts of the network. For the cable MSO with a lot of unique real estate and right of way investments, this is a powerful competitive advantage.

## 5. A Converged Edge Architecture

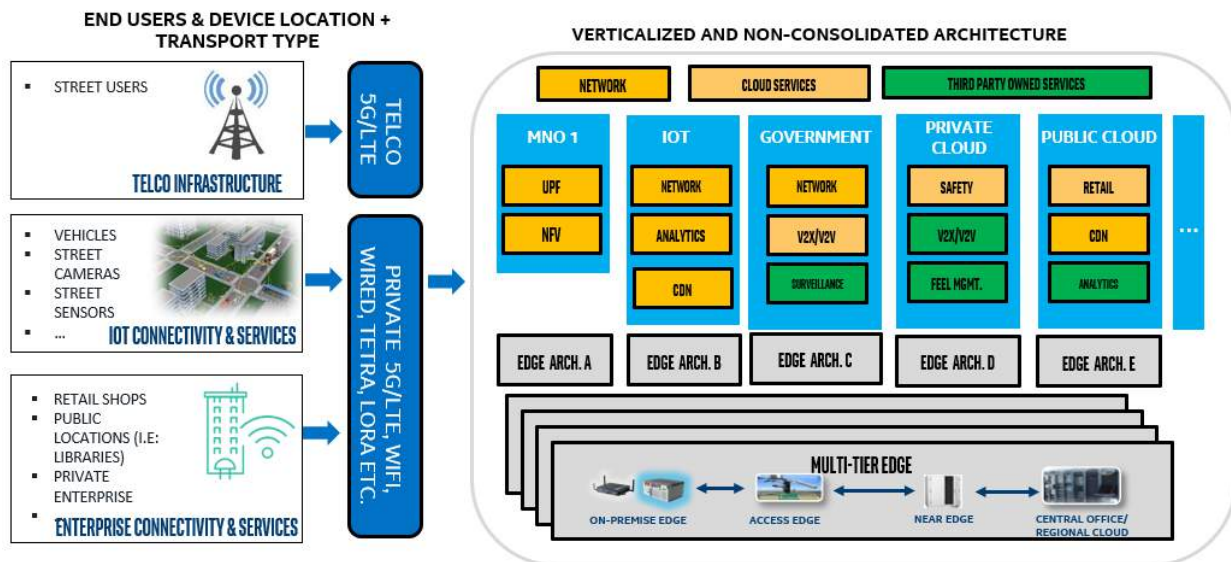
Once an Edge architect has determined which locations can support Edge platforms, there is still the choice of which equipment and software to deploy to execute the desired On-Prem or Network Edge models and the subsequent business arrangements discussed earlier. If a CoSP seeks to cost-effectively develop and scale more than one Access technology and Service infrastructure for the Edge, there must be some common platform to run everything.

The idea of minimizing the number of platforms comes under the banner of “convergence.” CableLabs identified a Converged Network Architecture Framework that defines the different types of convergence that apply to the 10G network:

- Access Convergence
- Transport Convergence
- Platform Convergence
- Core Convergence
- Operations Convergence

The “platform” or Edge platform that has been discussed in this paper is a concept that, in the ideal case, spans across all these domains. That is, how does one design Network Edge equipment and software infrastructure to enable a seamless user experience across all Access types, have a limited set of common hardware, take the best elements of the cloud world, consolidate management of Services, and ease operations with telemetry and automation? On top of that, the platform needs to be a carrier-class solution, computationally lightweight and efficient, high performance, and have facilities for optimized life cycle management.

In the past, it was not possible to set up a common hardware and software infrastructure to meet the needs of all areas of a network, support multiple Access technologies, and offer multiple Services in more than one domain. Figure 14 shows how technical solutions associated with various business owners were set up in silos, and thus they were developed, deployed, and managed independently. In this scenario, each solution is bespoke, and in a world where only one or two solutions was needed at a time, it was enough and still cost-effective to develop such independent systems and institutional expertise.



**Figure 14 - Access and Service Infrastructure in Silos**

However, the future 10G network needs to accommodate many wired and wireless Access technologies, and a host of new Services; and consequently, the old way of architecting networks is not going to work from economic or operational perspectives. What ameliorates this challenge is the rise of virtualization, containerization, and cloudification, new standards for power management, telemetry, and slicing the network for different Service tiers, and the resulting ability to converge a multitude of workloads on the same COTS equipment (i.e., standards-based servers and switches). But there still needs to be a

framework or a set of building blocks to put everything together in a way that scales across all the locations in the network, from the large data centers to the Nodes.

There are already several industry efforts towards software platforms at the Edge, including OpenNESS, CNTT Edge, Project Adrenaline, OpenVINO, and Open Visual Cloud. Some of these efforts have a relatively broad scope, for example, to move software infrastructure for the cloud to what is presumed to be a scaled-down platform for the Edge. Other efforts focus only on addressing the needs of specific domains (i.e., visual processing). However, to really touch on all aspects of convergence, there needs to be a view and a framework that can aid both hardware and software design using scalable building blocks so solutions can be deployed in any part of the Edge network.

With that in mind, researchers at Intel, led by Francesc Guim, Principal Engineer, and Timothy Verrall, Senior Principal Engineer, for the Intel Edge Architecture Group,<sup>7</sup> have been contributing to the development of a framework called the Converged Edge Architecture. The objective of this effort is to unify and converge Access, IoT, and other workloads on standards-based hardware and software. This Converged Edge Architecture is not a specific piece of hardware or software, but rather it is a set of building blocks to create a “Plug and Play” Edge platform that is relevant to the specific requirements of CoSPs for any location and for any set of functions.

The base of the Converged Edge Architecture framework assumes the hardware can be constructed using components that provide common features important for Edge deployments across the full range of performance needs and power constraints found in an operator’s network. For an Edge platform, these base features need to include easy and performant virtualization, large software support across many vertical domains, extendibility through accelerators, strong security capabilities, and moving forward, functions for real-time machine learning algorithms. Silicon supporting x86 architecture satisfy all these criteria today, but in the future, other options may emerge.

Possible hardware configurations for a Converged Edge Architecture-based solution are:

- Headend: a rack of x86-based 1RU or 2RU servers with a programmable top of rack switch
- Node: a small form factor x86 server for a standardized GAP enclosure
- Outdoor uCPE: a small form factor x86 server in a custom ruggedized enclosure

In all these cases, a common software management infrastructure identifies the compute, storage, and networking capabilities and connectivity of each location and orchestrates Access functions and Services according to defined service level agreements. Figure 15 shows the high-level design for Converged Edge Architecture software infrastructure. It consists of a sub-infrastructure to host the data plane functions for Access technologies, sub-infrastructure to host Services, a transport/switching infrastructure to move data from hardware (e.g., NICs) to the dataplane or Services or between the dataplane and any of the Services, and an infrastructure for coordinated deployment, orchestration, and management of all elements therein.

---

<sup>7</sup> Francesc Guim, Principal Engineer, Timothy Verrall, Senior Principal Engineer Intel Edge Architecture Group, 2020.

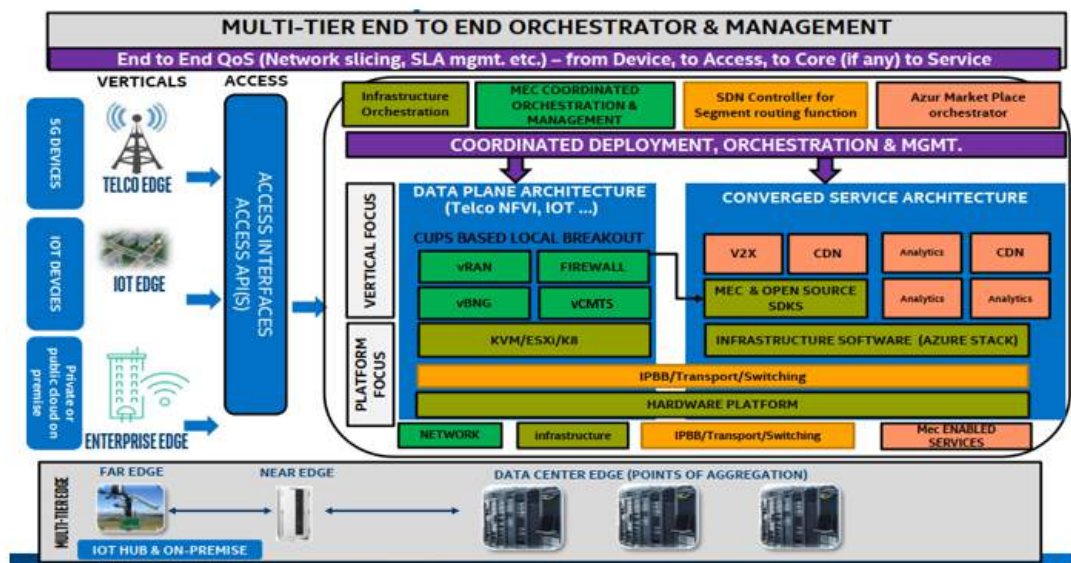
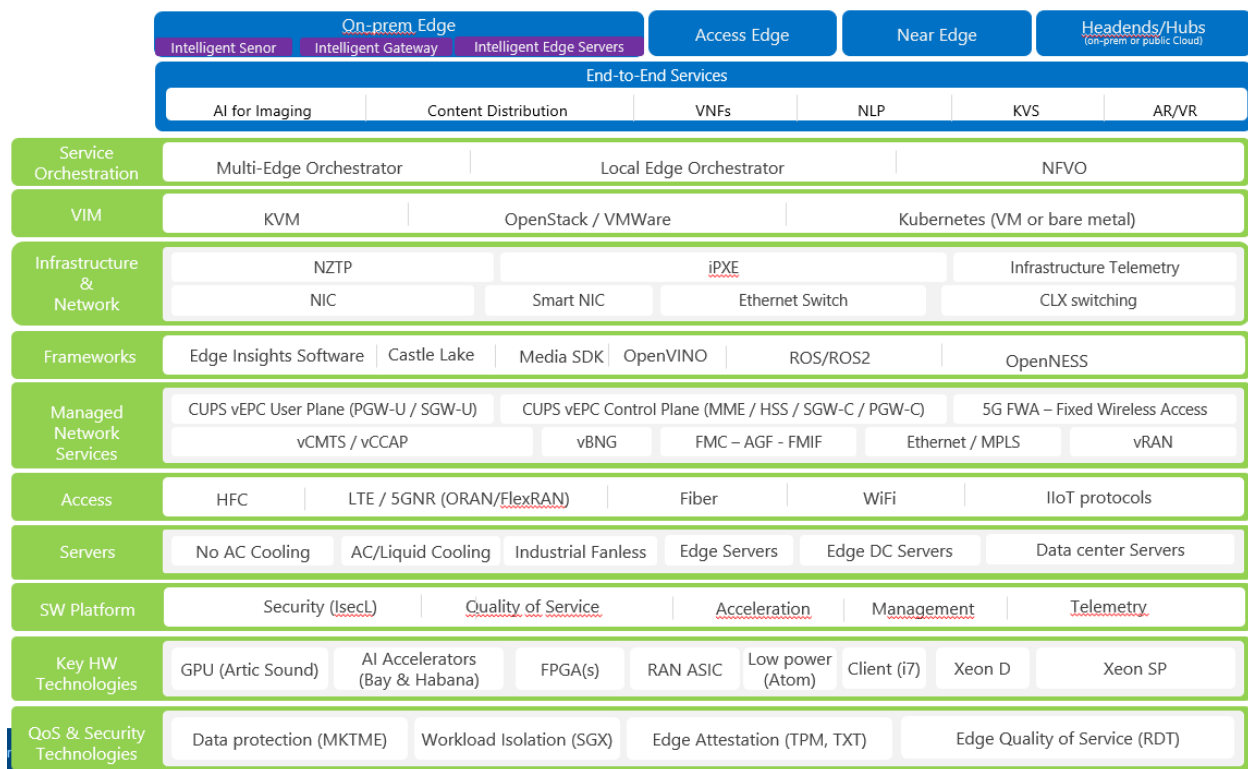


Figure 15 - Converged Edge Architecture

Figure 16 shows the next level down of granularity for types of software components the Converged Edge Architecture framework considers within the previously mentioned sub-infrastructures. Also shown is a sampling of candidate options from the market, including some of the familiar names in the industry. A solution would include at least one option from each row, though it would be common to include multiple elements. For example, a Smart Cities platform may have both Wi-Fi and HFC Access technologies, supported along with two or more frameworks to perform local video analytics and execute action plans.





**Figure 16 - Mix and Match for an Edge Platform**

Work on Converged Edge Architecture continues as real-life solutions for the IoT and Network Edge solutions are developed, new hardware and software elements are deployed, and lessons are learned in real-world conditions.<sup>8</sup> The goal is to provide a framework or template that can be used by OEMs, ODMs, TEMs, and ISVs to develop scalable, flexible, Edge platform solutions for wherever they may be needed across the Edge.

## 6. Considerations for Designing an Edge

The MSO network of the 10G era is Multi-Access, allows next-gen Services in a wide range of performance tiers, and has the flexibility to deploy capabilities where they are needed in the network to deliver on business goals. Some of these new Services require lower latencies or tighter controls around data sovereignty, which leads us to look at the capabilities of facilities closest to the end users – infrastructure collectively known as “the Edge”.

### 6.1. A Summary of the Options

This Edge, of course, is not necessarily constrained to being one location, one platform, or one type of business arrangement. At the top level, an operator can plan for an On Prem Edge or Network Edge – or both. The On Prem Edge implies that equipment at the customer site, like a uCPE, will run one or more Services locally while being managed and controlled centrally in the network. The Network Edge is

<sup>8</sup> “Converged Edge Reference Architecture (CERA) for On-Premise/Outdoor”, <https://builders.intel.com/docs/networkbuilders/converged-edge-reference-architecture-cera-for-on-premise-outdoor.pdf>

based on hosting Access and Services from equipment owned and operated by the CoSP, partner CSPs, or other 3<sup>rd</sup> parties in various arrangements of a Network Edge platform architecture:

- CoSP + CSP Co-location
- CoSP Led + CoSP/CSP Services
- CSP Led
- CoSP/CSP Aggregator

The platform architecture for each of these is made of a limited set of hardware (i.e. if virtualization and programmable components are involved) and software infrastructure to manage different elements of the Edge solution as well as the interface(s) to outside network elements. The details are in an earlier section, but the difference between the architecture types really boils down to the questions of:

- Who owns the real estate?
- Who owns the physical equipment?
- Who owns the software infrastructure(s)?
- Who owns the data/customer?

Generally speaking, ownership gives more opportunity for monetization, but it also means the owning organization needs more institutional knowledge in the domain at hand to make sure the technology delivers the desired results.

The MSO network, with years of developing and deploying an HFC plant, is in a unique position to “own the real estate”, that is, where to host Edge Platform equipment to provide the best latency and data locality profiles for new Services. In fact, intelligence and compute capabilities are coming to all places in the network – even to Nodes and other locations in the outside plant. Along these lines, the MSO/CoSP also has the expertise to design and deploy whatever Access elements required for the Edge solutions, although partnerships may be involved when new Access technologies, like a 5G wireless Service, are added to the network.

Where it gets more interesting – and where there is a lot of innovation/experimentation happening – is in answering the rest of the questions as they relate to providing Services. It seems the default behavior would be to bring in a CSP or 3<sup>rd</sup> party to host and manage their own Services over the last mile broadband connection being provided by the CoSP. This is simply because the Cloud technologies used in this case might be outside of the core competencies of the CoSP, so partnering with a CSP is the most straightforward way to monetize the aforementioned real estate advantages.

In the CoSP + CSP Co-location model, the CSP physically houses their equipment in the same location as the Edge location of the CoSP, be it a Headend, Hub, or Node. The CSP Led model is similar, except that the CSP equipment resides in a point of presence near-to, but outside, of the CoSP Edge location. In these cases, the CSP and/or 3<sup>rd</sup> parties making use of CSP resources “own the physical equipment... software infrastructure... data” and therefore the customers for the Services being sold. They could even host Services that the CoSP wants to provide like localized SD-WAN offerings for small and medium businesses.

CoSP’s wanting to follow the CoSP Led + CoSP/CSP Services model and offer their own Edge infrastructure to host Services – i.e. to own the equipment and software infrastructure – will have to develop or hire their own expertise in Cloud technologies. This might seem like a difficult and far-out proposition to those used to single function appliances in their network, but as Access workloads get virtualized (ex. vCCAP, vBNG, vRAN, etc.) the technologies to manage and deploy both Access and

Services are starting to converge. The final alternative is to leave it to an aggregator with a local point of presence host the equipment and software infrastructure and sell broadband and, potentially, Services through them (along with perhaps similar offerings from competitors).

Again, a key competitive advantage for the MSO is that it has invested in and has rights of way for Edge-friendly locations to host Edge platforms for either itself or for a CSP. Even though these locations range in physical space available, the amount of power that can be delivered, and other environmental constraints, with virtualization, telemetric capabilities, facilities for remote security, and software infrastructure for managing such distributed computing elements, it is possible to consolidate the number of platforms into the minimum possible. Less disparate architectures and technological domains means a better total cost of ownership across the network.

To this end, the Converged Edge Architecture project is an effort to help OEMs, ODMs, TEMs, and ISVs mix and match common COTS hardware and open source software elements through a common framework to construct a scalable and flexible platform that matches the performance, lifecycle, and form factor requirements for any Edge location. Commercial examples of industrial On-Prem, 5G vRAN, and other solutions that made use of the Converged Edge Architecture framework are starting to emerge and providing the industry with proof points of the benefit leveraging standards, COTS hardware, and reference software architectures for faster time to market and lower total costs of ownership.

## **6.2. Asking the Right Questions**

Admittedly, it will take more than a checklist to consider all the options above and architect an Edge or Edges in a given network. But a few key decisions that will drive the planning and architecture are:

- What type of Services do you want to offer and what requirements do they have on the network?
- What business models / partnerships do you want to support – who owns what?
- Where are you willing to deploy equipment / functions / infrastructure?
- What equipment and software infrastructure can be consolidated across the network?
- Who is going to own the various parts of the Edge solution in the organization?

The last question about organizational ownership may be the hardest as “the Edge” crosses what were typically separate domains – multiple Access technologies, Enterprise Services, Data Center and Cloud resources, etc. But it is because of this breadth that it is clear the Edge is important for MSOs to grow their businesses; whatever options are chosen. In fact, while difficult, thinking about a grand Edge strategy may be the chance for an operator to re-think existing silos, align on the latest technological innovations, and be a distinguishing factor against the competition.

All that said, practical realities to leverage existing systems, skillsets, and business relationships will make this an iterative process. That is, it is good to develop the ultimate end-state for your network and the organizations supporting it, but you may have to accomplish the transformation in stages. For example, would implementing a CSP Led model first allow you to get into the market and show the value of your CoSP Edge locations while you start developing Cloud expertise internally and drive the Ecosystem to be able to consolidate Access and Services later onto the same servers and switches?

Regardless of the path, the time is now to get started on this journey! MSOs can leverage their unique infrastructure and real-estate investments to provide improved network visibility, performance, control, flexibility, and agility with a distributed compute architecture all the way to the Edge – wherever it may be!



# Abbreviations

|        |                                                 |
|--------|-------------------------------------------------|
| API    | application programming interface               |
| CERA   | Converged Edge Reference Architecture           |
| CPE    | customer premise equipment                      |
| CSP    | cloud service provider                          |
| CoSP   | communications service provider                 |
| CNF    | cloud native function                           |
| DAA    | Distributed Access Architecture                 |
| GAP    | Generic Access Platform                         |
| ISV    | independent software vendor                     |
| KPI    | key performance indicator                       |
| MSO    | multi-service operator                          |
| NFVI   | network functions virtualization infrastructure |
| ODM    | original design manufacturer                    |
| OEM    | original equipment manufacturer                 |
| OTT    | over-the-top                                    |
| PAAS   | platform-as-a-service                           |
| SD-WAN | Software-defined wide area network              |
| SCTE   | Society of Cable Telecommunications Engineers   |
| TEM    | telecommunications equipment manufacturer       |
| uCPE   | universal customer premise equipment            |
| VNF    | virtual network function                        |

# Bibliography & References

Gartner, “Edge computing promises near real-time insights and facilitates localized actions.”; <https://www.gartner.com/smarterwithgartner/what-edge-computing-means-for-infrastructure-and-operations-leaders>

Xu et. al., “Edge Intelligence: Architectures, Challenges, and Applications”, <https://arxiv.org/abs/2003.12172>

FCC, “2016 Measuring Broadband America Fixed Broadband Report”; <https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-fixed-broadband-report-2016>

Fierce Wireless, “Verizon, AWS bring 5G MEC to Boston, Bay Area”; <https://www.fiercewireless.com/operators/verizon-aws-bring-5g-mec-to-boston-bay-area>

Mann, “Coronavirus Cleaves SD-WAN Revenue Growth”; <https://www.sdxcentral.com/articles/news/coronavirus-cleaves-sd-wan-revenue-growth/2020/06/>

RTTNews.com, “Microsoft To Use Telefonica Infrastructure For Datacenter Region In Spain” ; <https://www.nasdaq.com/articles/microsoft-to-use-telefonica-infrastructure-for-datacenter-region-in-spain-2020-02-26>

Robuck, “Cox targets the edge for the next evolution of network performance and security”, <https://www.fiercetelecom.com/telecom/cox-targets-edge-for-next-evolution-network-performance-and-security>

Dano, “SBA, American Tower double down on edge computing opportunity”; <https://www.lightreading.com/the-Edge/sba-american-tower-double-down-on-Edge-computing-opportunity/d/d-id/762941>

Levensalor, Stuart, “The Modular, Virtualized Edge for the Cable Access Network”; <https://community.cablelabs.com/wiki/plugins/servlet/cablelabs/alfresco/download?id=2c46cef2-af44-47be-bdd4-98a948cbc60d>

SCTE Generic Access Platform; <https://www.scte.org/generic-access-platform/>

Open Network Edge Services Software; <https://www.openness.org/>

CableLabs Project Adrenaline; <https://www.cablelabs.com/edge-adrenaline-boost-kubernetes-orchestrate-fpgas-gpu>

OpenVINO toolkit; <https://docs.openvinotoolkit.org/>

Open Visual Cloud; <https://01.org/openvisualcloud>

“Converged Edge Reference Architecture (CERA) for On-Premise/Outdoor”, <https://builders.intel.com/docs/networkbuilders/converged-edge-reference-architecture-cera-for-on-premise-outdoor.pdf>

# **Covid-19 Learnings**

## **All Roads Lead To DOCSIS® 4.0 Technology**

A Technical Paper prepared for SCTE•ISBE by

**Jeff Finkelstein**

Chief Access Scientist  
Cox Communications  
6305-B Peachtree Dunwoody Rd, Atlanta, GA 30328  
+1 404.269.5955  
Jeff.Finkelstein@cox.com

**Tom Cloonan**

CTO – Network Solutions  
CommScope  
2400 Ogden Ave, Suite 180, Lisle, IL 60532  
+1 630.281.3050  
Tom.Cloonan@commscope.com

**Doug Jones**

Principal Architect  
CableLabs®  
858 Coal Creek Circle  
+1 303.661.9100  
D.Jones@cablelabs.com

# Table of Contents

| <b>Title</b>                                         | <b>Page Number</b> |
|------------------------------------------------------|--------------------|
| 1. Introduction.....                                 | 3                  |
| 2. Background.....                                   | 3                  |
| 3. Traffic Engineering Model .....                   | 4                  |
| 4. Early Learnings .....                             | 11                 |
| 5. More Upstream Means a Few Technology Changes..... | 16                 |
| 6. Current Activities.....                           | 16                 |
| 7. Longer-Term Activities .....                      | 17                 |
| 8. Conclusion .....                                  | 18                 |
| Abbreviations.....                                   | 19                 |
| Bibliography & References .....                      | 19                 |

## List of Figures

| <b>Title</b>                                                                                      | <b>Page Number</b> |
|---------------------------------------------------------------------------------------------------|--------------------|
| Figure 1 - Required Downstream Bandwidth Capacity vs. Time for “Typical MSO” (pre-COVID-19).....  | 6                  |
| Figure 2 - Required Upstream Bandwidth Capacity vs. Time for “Typical MSO” (pre-COVID-19) .....   | 7                  |
| Figure 3 – Example Downstream Bandwidth Changes as “Stay-At-Home” was Rolled Out.....             | 8                  |
| Figure 4 – Example Upstream Bandwidth Changes as “Stay-At-Home” was Rolled Out .....              | 8                  |
| Figure 5 – Required Downstream Bandwidth Capacity vs. Time for “Typical MSO” (post-COVID-19)..... | 10                 |
| Figure 6 – Required Upstream Bandwidth Capacity vs. Time for “Typical MSO” (post-COVID-19).....   | 10                 |
| Figure 7 – Large Guard Bands Between Upstream Channels .....                                      | 12                 |
| Figure 8 – Increasing Channel Width To Increase Upstream Capacity .....                           | 13                 |
| Figure 9 – Maximizing Upstream Capacity With Wide Upstream Channels and Small Guard Bands.....    | 13                 |
| Figure 10 – Adding Additional Upstream Channels .....                                             | 13                 |
| Figure 11 – DOCSIS 3.1 Upstream Spectrum Options .....                                            | 14                 |
| Figure 12 – DOCSIS 4.0 Upstream Spectrum Options .....                                            | 15                 |

## List of Tables

| <b>Title</b>                                | <b>Page Number</b> |
|---------------------------------------------|--------------------|
| Table 1 – Upstream Capacity Comparison..... | 15                 |

## 1. Introduction

My most used quote for 2020 so far can be summed up in one of my Bubbie's (Yiddish for grandmother) favorite sayings, which is as follows: *Der Mentsh tracht un Gott lacht*. Which is roughly translated as "Man plans and God laughs".

Despite our plan of intents, plan of records, short-range plans, and long-range plans, the World was caught by surprise with the current events impacting us. We have seen our carefully laid plans cast aside as we have all scrambled to meet user demands resulting from mass isolation to curb the spread of an unseen enemy. For many of us we have experienced 18-24 months of utilization growth in a 3-month window from March to June 2020. Thankfully we are fortunate to have deployed spare capacity and use a data transmission protocol that has handled the new demands very well.

With ubiquitous stay-at-home orders being issued throughout the world in the March and April timeframe, a new era and social experiment had begun. It has created conditions that had never before been witnessed on the planet in the modern era; it would stress many areas of society, including medical resources, human relations, business operations, educational institutions, entertainment and sports, politics and elections, food distribution, travel, supply chains, and much more.

There has been and will be much written about causes and effects of the current World condition. In this paper we will examine learnings we have gathered so far during this unparalleled journey focusing on the cable network and DOCSIS technology in particular.

## 2. Background

A key segment of society that was undoubtedly surprised and, in some cases, strained by these rapid changes was the high-speed data broadband infrastructure that makes up an important portion of the Internet. With more people staying at home and working at home than ever before, the COVID-19 pandemic introduced rapid and profound changes to the high-speed data infrastructure.

In the DOCSIS networks operated by MSOs, these changes were felt almost instantaneously in both the upstream and downstream data paths within HFC plants. It is accurate to state that no traffic engineers within the Cable Industry (or any Industry for that matter) had ever anticipated the potential need for stressful traffic engineering models simulating bandwidth consumptions that might result from the arrival of a worldwide epidemic that would drive a majority of the subscriber base into their homes for 24 hours a day and for many months at a time.

The surprising and unexpected arrival of these incredible conditions created an intense traffic load, and this load could have easily overloaded the DOCSIS network and rendered the infrastructure (and the Internet) useless at a time when "work-at-home" and the ability to connect to others on-line became a necessity for society.

However, the DOCSIS infrastructure and the Internet held up quite well to these unexpected stresses. Granted, there were some growing pains and problematic issues that required emergency actions for correction, but many of these issues were rapidly resolved with CMTS/CM configuration changes and/or channel augmentations. The Cable Industry was able to successfully respond to this sudden onslaught of traffic. But the issues that were encountered by this sudden bandwidth surge were likely exposing an Achilles Heel that will need to be addressed in the near future.

### 3. Traffic Engineering Model

It may be instructive to explore some typical traffic engineering models that may have been being used by MSOs in February 2020 prior to the arrival of the COVID-19 pandemic. Within these models, the authors will make use of a “QoE-based Traffic Engineering Formula” defined in [1], which is given by:

$$\text{Required Bandwidth Capacity} = N_{\text{sub}} * T_{\text{avg}} + K * T_{\text{max}} \quad (1)$$

where:

- $N_{\text{sub}}$  is the number of subscribers in the Service Group
- $T_{\text{avg}}$  is the Average Bandwidth Consumption of a single subscriber (during the busy hour)
- $T_{\text{max}}$  is the maximum Bandwidth offered in the Service Level Agreements
- $K$  is the QoE coefficient, which is a value typically between 0.8 and 1.5 (larger values yield better Quality of Experience);  $K=1.2$  is the value used in this paper.

Let us consider a “typical” HFC network in February of 2020 that might have been architected for a  $N_{\text{sub}}=250$ -subscriber Service Group with the following assumptions:

- Downstream  $T_{\text{max}} = 1000$  Mbps
- Downstream  $T_{\text{avg}} = 2.5$  Mbps
- Downstream  $T_{\text{avg}}$  CAGR of 26%
- Upstream  $T_{\text{max}} = 30$  Mbps
- Upstream  $T_{\text{avg}} = 0.175$  Mbps
- Upstream  $T_{\text{avg}}$  CAGR of 26%

Using these assumptions, the following Traffic Engineering rules are observed:


- “Typical” High-Speed Data Downstream Traffic Engineering in February of 2020
  - 32 x 6 MHz 256QAM SC-QAM Channels @ 36 Mbps per channel
    - $32 * (36 \text{ Mbps}) = 1152 \text{ Mbps Capacity}$
  - 1 x 192 MHz 1024QAM OFDM Channel @ usable spectral efficiency of 10 bps/Hz
    - $(192 \text{ MHz}) * (10 \text{ bps/Hz}) = 1920 \text{ Mbps Capacity}$
  - Total Downstream Bandwidth Capacity

- $1152 \text{ Mbps} + 1920 \text{ Mbps} = 3072 \text{ Mbps}$
- Busy-hour Downstream Average Bandwidth Consumption per Subscriber = 2.5 Mbps
- Number of Subscribers in a Typical Service Group = 250 Subscribers
- Average Busy-hour Downstream Bandwidth Consumption ( $T_{\text{avg}}$ ) in a Typical Service Group
  - $250 * (2.5 \text{ Mbps}) = 625 \text{ Mbps}$
- “Headroom” Downstream Bandwidth (to support traffic bursts) in a Typical Service Group
  - $3072 \text{ Mbps} - 625 \text{ Mbps} = 2447 \text{ Mbps}$
- Maximum Supportable Downstream SLA Throughput ( $T_{\text{max}}$ ) calculated using the “QoE-based Traffic Engineering formula with a K of 1.2” = Headroom/K =  $2447 \text{ Mbps} / 1.2 = 2040 \text{ Mbps}$
- Maximum Downstream Life-span of existing HFC Plant (without changes) using the “QoE-based Traffic Engineering formula with a K of 1.2 and Downstream  $T_{\text{avg}}$  CAGR of 26%” and  $T_{\text{max}}$  of 1000 Mbps”
  - $\log[(3072 \text{ Mbps} - 1.2 * 1000 \text{ Mbps}) / 625 \text{ Mbps}] / \log[1.26] = 4.75 \text{ years}$
- “Typical” High-Speed Data Upstream Traffic Engineering in February of 2020
  - 3 x 6.4 MHz 64QAM ATDMA Channels @ 25 Mbps per channel
    - $3 * (25 \text{ Mbps}) = 75 \text{ Mbps Capacity}$
  - 1 x 3.2 MHz 64QAM ATDMA Channels @ 12.5 Mbps per channel
  - 1 x 9.6 MHz OFDMA Channels @ usable spectral efficiency of 8 bps/Hz
    - $(9.6 \text{ MHz}) * (8 \text{ bps/Hz}) = 76.8 \text{ Mbps Capacity}$
  - Total Upstream Bandwidth Capacity
    - $75 \text{ Mbps} + 12.5 \text{ Mbps} + 76.8 \text{ Mbps} = 164.3 \text{ Mbps}$
  - Busy-hour Upstream Average Bandwidth Consumption per Subscriber = 0.175 Mbps
  - Number of Subscribers in a Typical Service Group = 250 Subscribers
  - Average Busy-hour Upstream Bandwidth Consumption ( $T_{\text{avg}}$ ) in a Typical Service Group
    - $250 * (0.175 \text{ Mbps}) = 43.75 \text{ Mbps}$
  - “Headroom” Upstream Bandwidth (to support traffic bursts) in a Typical Service Group

- $164.3 \text{ Mbps} - 43.75 \text{ Mbps} = 120.55 \text{ Mbps}$
- Maximum Supportable Upstream SLA Throughput ( $T_{\max}$ ) calculated using the “QoE-based Traffic Engineering formula with a K of 1.2”
  - $\text{Headroom}/K = 120.55 \text{ Mbps}/1.2 = 100 \text{ Mbps}$
- Maximum Upstream Life-span of existing HFC Plant (without changes) using the “QoE-based Traffic Engineering formula with a K of 1.2 and Upstream  $T_{\text{avg}}$  CAGR of 26%” and  $T_{\max}$  of 30 Mbps”
  - $\log[(164.3 \text{ Mbps} - 1.2 \times 30 \text{ Mbps})/43.75 \text{ Mbps}] / \log[1.26] = 4.6 \text{ years}$


Note: The illustrative numbers above do not represent any particular MSO. Different MSOs will have different Traffic Engineering numbers and different Bandwidth Capacity numbers than those shown above.

As can be seen in the example scenarios above, this “typical” MSO in February would have had enough “headroom” in their Downstream Bandwidth Capacity to permit them to operate for another 4.75 years prior to having to consider any type of plant modifications to accommodate the expected Downstream  $T_{\text{avg}}$  Bandwidth growth rate (resulting from a predicted 26% CAGR that might have existed in February). In addition, this “typical” MSO in February would also have had enough headroom in their Upstream Bandwidth Capacity to permit them to operate for another 4.5 years prior to having to consider any plant modifications to accommodate the expected Upstream  $T_{\text{avg}}$  Bandwidth growth rate (resulting from a predicted 26% CAGR that might have existed in February). These lifespans are shown by the Bandwidth Requirements curves in Figure 1 and Figure 2 which are based on the QoE-based Traffic Engineering formula of Eq. (1) above.



**Figure 1 - Required Downstream Bandwidth Capacity vs. Time for “Typical MSO” (pre-COVID-19)**





**Figure 2 - Required Upstream Bandwidth Capacity vs. Time for “Typical MSO” (pre-COVID-19)**

Even without COVID-19, this “typical” MSO would have likely “run out of gas” in their Upstream during the fourth quarter of 2024, and it would have forced the subscribers to live with sub-standard QoE for a while or it would have forced the MSO to utilize one of several options to modify their plant to extend the capacity and life-span of their DOCSIS Upstream services. Many of those possible options for expanding the Available Bandwidth Capacity will be outlined in sections below.

Now let us consider the profound bandwidth usage changes that occurred in the mid-to-late March timeframe in most MSOs and let us explore their impact on the curves and DOCSIS lifespans of Figure 1 and Figure 2.

The bandwidth activities of post-COVID-19 “stay-at-home” subscribers are very different from the activities of subscribers prior to the arrival of COVID-19; in both the upstream and downstream direction, and during the day and at night. In a matter of a few days while “stay-at-home” was becoming the norm, the bandwidth demands increased more rapidly than ever seen before.

Reported Downstream Bandwidth consumption ( $T_{avg}$ ) during the evening busy-hour timeframe increased by 12-25% at different MSOs. Upstream Bandwidth consumption ( $T_{avg}$ ) during the evening busy-hour timeframe increased by 20-50% at different MSOs. Most of this night-time bandwidth growth was driven by a combination of more video streaming, video-conference work meetings, social networking traffic, and gaming traffic (both playing and viewing). In the Upstream, it also included an increase in TCP ACKs associated with the Downstream packets.


Reported Downstream Bandwidth consumption ( $T_{avg}$ ) during the mid-day timeframe increased by 26-86% at different MSOs and reported Upstream Bandwidth consumption ( $T_{avg}$ ) during the mid-day timeframe increased by 30-150% at different MSOs. Most of this day-time bandwidth growth was driven by a combination of video-conference work meetings, videoconferencing for remote education, social networking traffic, and gaming traffic (both playing and viewing). In the Upstream, it also included an increase in TCP ACKs associated with the Downstream packets.

Examples changes with US operators from the NCTA web page are shown in Figure 3 and Figure 4.

## National Downstream Peak Growth

### Observed Increase in Peak Consumer Usage

Overall Change in Pre-COVID Internet Usage Since Early March Compared to the Weekly Usage Change



**Figure 3 – Example Downstream Bandwidth Changes as “Stay-At-Home” was Rolled Out**

## National Upstream Peak Growth

### Observed Increase in Peak Consumer Usage

Overall Change in Pre-COVID Internet Usage Since Early March Compared to the Weekly Usage Change



**Figure 4 – Example Upstream Bandwidth Changes as “Stay-At-Home” was Rolled Out**

In the figures above, the busy-hour bandwidth (at night) still exceeds the mid-day bandwidth, so the busy-hour bandwidth characteristics will still stress the DOCSIS system and define the ultimate Bandwidth Capacity Requirements (as they did in the past prior to COVID-19).

For our “typical” MSO scenario (post-COVID-19), let us assume that they experienced the following  $T_{avg}$  BW growth spurts on their networks when “stay-at-home” rules were initiated:

- Downstream  $T_{avg}$  Bandwidth step function increase in 2Q2020 = 25%
- Upstream  $T_{avg}$  Bandwidth step function increase in 2Q2020 = 50%


It is assumed (and hoped) that there will come a time in different localities when “stay-at-home” orders will begin to be relaxed. At that point in time, it is expected that some workers will begin to return to work at the office. However, the success of the “work-at-home” model during 2020 may push many companies to consider allowing more of their employees to continue working at home even after the “stay-at-home” orders are relaxed. We will create models for our “typical” MSO which assume that 25% of the “stay-at-home” workers will return to the office in 3Q’2020 and another 25% of the “stay-at-home” workers will return to the office in 1Q’2021. Note: These percentages may be accomplished by rotating different employees into the office at different times to reduce the density of employees in the office and reduce the likelihood of COVID-19 transmissions).

After the two phases of employees returning to work, we are assuming that 50% of the “stay-at-home” workers will continue to work at home moving into the future. As a result, it will leave the Bandwidth requirements higher in the long-term; even after COVID-19 has returned the world back to the “new normal.” These “new normal” Bandwidth requirements will thus be higher than they would have been had COVID-19 never existed.


The resultant modified Bandwidth Requirements curves (post-COVID-19) are shown in Figure 5 and Figure 6. Within these figures, we have assumed that the network conditions have not changed (other than the step functions resulting from COVID-19).

We are still assuming an  $N_{sub}=250$ -subscriber Service Group with the following assumptions:

- Downstream  $T_{max} = 1000$  Mbps
- Downstream  $T_{avg} = 2.5$  Mbps
- Downstream  $T_{avg}$  CAGR of 26%
- Upstream  $T_{max} = 30$  Mbps
- Upstream  $T_{avg} = 0.175$  Mbps
- Upstream  $T_{avg}$  CAGR of 26%



**Figure 5 – Required Downstream Bandwidth Capacity vs. Time for “Typical MSO” (post-COVID-19)**



**Figure 6 – Required Upstream Bandwidth Capacity vs. Time for “Typical MSO” (post-COVID-19)**

Comparing both Figure 1 to Figure 5 and Figure 2 to Figure 6, it is clear that COVID-19 may play a large role in DOCSIS network evolution going forward. It may force MSOs to make network transitions sooner than they had originally planned - even if the Downstream  $T_{avg}$  CAGRs and Upstream  $T_{avg}$  CAGRs do not increase. The large bandwidth “step function” up resulting from “stay-at-home” actions and the smaller bandwidth “step functions” down resulting from limited “return-to-work” actions leads to these transitions.

It can be seen in Figure 5 that the impact of the COVID-19 step functions would force Downstream network modifications for this “typical” MSO to occur in 2Q2024 instead of 4Q2024. And it can be seen in Figure 6 that the impact of the COVID-19 step functions would force Upstream networks modification for this “typical” MSO to occur in 2Q2024 instead of 4Q2024

This “typical” MSO may not represent the situations at all MSOs. Many MSOs have smaller Service Groups with less than 250 subscribers. And many MSOs also have allotted more Bandwidth Capacity to DOCSIS networks. For example, European MSOs typically have 65 MHz of spectrum dedicated to the Upstream, which gives them much more Upstream Bandwidth Capacity and a longer lifespan on their Upstream. And some MSOs throughout the world have already begun to transition to 85 MHz Mid-Splits. As a result of this additional “headroom” (in the form of extra Bandwidth Capacity), many of these MSOs were able to easily absorb the bandwidth surge that occurred during the “stay-at-home” orders, and they therefore saw little or no problems on their Upstream DOCSIS systems when COVID-19 hit.

## 4. Early Learnings

DOCSIS technology has handled the dramatic increase in usage very well. Both work-at-home and schooling-at-home created a large step-function in demand where operators have seen 18 months of growth in the first 3 months of isolation. In fact, with the effort to flatten the curve everything became focused in the home including work, school, entertainment, gaming, and new areas including exercise, tele-health, remote-library (book) access, video calls, etc. Just about everything migrated to the home broadband connection.

Why did DOCSIS broadband hold up so well? Starting from the ground up, the coaxial cable used for DOCSIS broadband is a superior medium from which more capacity keeps being mined. Coaxial cable supports much higher speeds than the twisted pairs used for telephone DSL service and even early fiber optic point-to-multipoint protocols. Starting with the DOCSIS 3.1 specifications in 2016, the coaxial cable supported higher speeds than a fiber optic network running GPON technology (gigabit passive optical network). The recent DOCSIS 4.0 specifications position cable-based broadband to rival 10 gigabit PON technologies using the existing coaxial cables in place today. This bandwidth can be realized in that at one point the coaxial cable carried hundreds of cable TV channels, which over the past decade have been migrated to broadband and entertainment migrating to IPTV.

Beyond the cable, DOCSIS technology has been optimized for orderly traffic flow. It has evolved to become a point-to-multipoint network that supports large amounts of capacity; and the network efficiently schedules the use of that capacity among users to meet service demand and provide quality of experience. The work first started in 1996 with collaboration among suppliers and operators to unify cable broadband technologies. Back then there were over a dozen proprietary systems which created a fragmented market and more importantly prohibited the retail availability of cable modems. The original goal was to unify around a single specification that brought together best-of-breed technologies. Without going into details, though very interesting, the effort was successful.

A huge benefit to unifying around a single set of specifications was getting the supplier community to focus their efforts on making DOCSIS technology the best-in-class service it has become today. The technical details include the physical layer (PHY), the media access control (MAC) layer, the security layer (SEC), operating system support (OSSI), and more. Over the next 20 years the innovations continued including advancements in silicon for better digital signal processing allowing even higher speeds and larger capacity to be realized.

The result is coupling an outstanding medium, i.e. coaxial cable, and very focused technology development to create the DOCSIS cable broadband as we know it today, with gigabit services possible anywhere there is a cable system. The DOCSIS 4.0 specifications support multigigabit symmetric service which will be the service offerings of this decade as that technology becomes available.


Both capacity and usage are instrumented in the DOCSIS network. There is a deep knowledge of how the network is performing both from a capacity and performance perspective on a day-to-day basis. The goal is to stay ahead of growth and the DOCSIS network has multiple levers to pull to increase capacity. Operators are always increasing capacity by one of two methods: adding more spectrum or segmenting the network so available spectrum is applied to fewer users.

The focus of DOCSIS 4.0 technology is increasing upstream capacity to as much as 6 Gbps. The operator has flexible options for upstream capacity, and along the way the downstream can be increased to as much as 10 Gbps.

Speaking to the upstream in North America, the vast majority of the coaxial cable is operated with a low-split (42 MHz) and typically has less capacity than the downstream. Low-split is a hold-over from the days of analog television channels and cable-ready TV sets (both of which are rapidly going the way of the do-do bird, if not already there). Low-split allowed the old Channel 2 to be carried on the coaxial cable starting at 54 MHz. The old analog Channel 2 has generally been replaced by digital TV carriers and IPTV, so the reasons for low-split are becoming few and far between.

DOCSIS equipment is very resilient, and for low-split operation it would encourage operators to experiment with these recommendations to get more upstream capacity. Next we discuss some strategies.


Removing large guard-band between carriers as it is not needed. The DOCSIS specifications are written such that RF channels can be directly adjacent to each other. In discussions with operators during the time of COVID-19, we have seen situations where the upstream DOCSIS carriers have large guard-bands between them as shown in Figure 7 below.



**Figure 7 – Large Guard Bands Between Upstream Channels**


These spaces between any DOCSIS carriers (upstream or downstream) have been found to not be needed; the carriers may be directly adjacent to each other on each side with minimal guard-band. There is a potential to use as little as 100 KHz spacing between ATDMA carriers. However, as you place a carrier near the roll-off starting at 42 MHz it is best to leave around a 500 KHz guard-band from 41.5 MHz to 42 MHz to avoid non-linearities related to amplifier cascades.

Increasing modulation order to 64 QAM for the legacy upstream DOCSIS carriers. The DOCSIS QAM technology is quite resilient, and the recommendation is to increase carriers to 64 QAM modulation, widen the carriers to use available spectrum, and remove the guard-band between the carriers. As shown in Figure 8 below, the moves should be done stepwise to confirm proper operation.




**Figure 8 – Increasing Channel Width To Increase Upstream Capacity**

Figure 9 shows using all available spectrum, which could include widening all upstream carriers and abutting those carriers.



**Figure 9 – Maximizing Upstream Capacity With Wide Upstream Channels and Small Guard Bands**

As an additional step, try one or more additional carriers, up high or down low as shown in Figure 10. Operators have been successful at adding new carriers, which is a testament to maintaining the plant more diligently over the last decade. The DOCSIS technology has many options for optimizing channel layout and is not the same as even a decade ago. The digital transceivers have increased in sensitivity and capabilities to enable optionality not thought of even 5 years ago.



**Figure 10 – Adding Additional Upstream Channels**

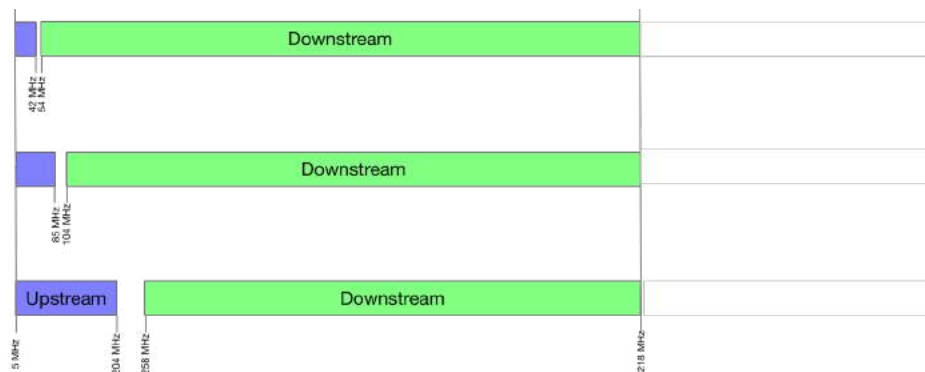
Homes passed decreased which lowers the effect of noise funneling at low frequencies. Cascades have shortened which lessens the impact of group delay close to the diplex filter cut-off frequency. Operators have been successfully running narrow carriers (typically with lower order modulation) both down to 10 MHz and closer to the diplex filter.

Next steps include:

- Allocating more spectrum to DOCSIS broadband, specifically in the upstream which has implications on the downstream
- Migrating to DOCSIS 3.1 technology that includes OFDMA in the upstream and new modulation orders up to 1024 QAM (Theoretically may be up to 4096 QAM)

The DOCSIS 3.1 specifications define three upstream spectrum options that include a top end of the spectrum at 1,218 MHz and are in use around the world, as shown in Figure 11

- Euro-split (65 MHz)
- Mid-split (85 MHz)
- High-split (204 MHz)




**Figure 11 – DOCSIS 3.1 Upstream Spectrum Options**

The DOCSIS 4.0 specifications include additional upstream spectrum options:

- 204 MHz
- 300 MHz
- 396 MHz
- 492 MHz
- 684 MHz



And the DOCSIS 4.0 specification includes operating the coax up to 1794 MHz, as shown in Figure 12.



**Figure 12 – DOCSIS 4.0 Upstream Spectrum Options**

Table 1 shows a comparison of the upstream capacities for these different upstream spectrum options shown in Figure 11 and Figure 12.

**Table 1 – Upstream Capacity Comparison**

| Upstream Frequency (MHz) | Increase in Spectrum | Aggregate Upstream Capacity |
|--------------------------|----------------------|-----------------------------|
| 42                       | Not applicable       | 100 Mbps                    |
| 85                       | 102%                 | 400 Mbps                    |
| 204                      | 385%                 | 1.4 Gbps                    |
| 300                      | 614%                 | 2.3 Gbps                    |
| 396                      | 842%                 | 3.1 Gbps                    |
| 492                      | 1,071%               | 4.0 Gbps                    |
| 684                      | 1,528%               | 5.7 Gbps                    |

The capacities above include 4 upstream SC-QAM channels and filling the rest of available spectrum with OFDMA channels (primarily 1024 QAM though your numbers may vary).

There are many options, but operators first have to wrap their heads around allocating more spectrum to the upstream. Which also may mean moving the top end of the downstream to allow for maintaining legacy video and data as increasing the return without changing the forward means less available forward spectrum.

## 5. More Upstream Means a Few Technology Changes

OFDMA technology, which includes high order modulations up to 1024 QAM technology, which is double the bit density of current upstream 64 QAM technology (10 bits per second per Hz as compared to 6 bits per second per Hz).

The nature of most internet network traffic has changed a lot over the past few decades. Many of the network protocols (at Layer 4 and higher) have become much more intelligent and more adaptive to network congestion. These algorithms use information on network packet loss and network packet delay to determine optimal throughputs that vary as a function of time. These algorithms (for the most part) tend to be very accommodating, backing off their traffic rates at the source whenever congestion is suspected anywhere within their traffic's path through the network. This is true of all of the Layer 4 TCP variants (ex: BIC, CUBIC) that have been developed and deployed in recent history. It is also true of the higher-layer protocols that throttle traffic for UDP-based applications (ex: QUIC). IP Video transport has evolved to now utilize many Adaptive Bit-Rate techniques at higher protocol layer to change the video resolution and the required video bandwidth capacity in response to network congestion. These types of adaptive bandwidth management techniques can be found in both MSO-managed IP video delivery services and OTT IP video delivery services. They can also be found in the video-conferencing solutions (ex: WebEx, Zoom, and Teams) that have become extremely popular for use by work-at-home employees. It was (in part) the ability of these protocols to throttle down their bandwidth usage during times of congestion that helped permit MSOs to survive the sudden bandwidth surge that occurred when the COVID-19 lock-downs first began. That throttling capability, coupled with the innovative scheduling algorithms in DOCSIS CMTSs, ensured that most users could get adequate QoE even in the presence of an incredibly high traffic load generated by their neighbors.

MSOs have taken great steps to increase the reliability of the cable broadband network. With the move to work-at home, network reliability takes front seat. No one wants network issues when on web calls most of the day, or students participating in classes from home. As part of its long-standing proactive network maintenance (PNM) project, CableLabs® has been developing technology and best practices to increase the reliability of the cable network. The SCTE took this effort up in 2017 by creating a PNM working group within its standards program, specifically in the Network Operations Subcommittee (NOS). PNM for the HFC network has taken advantage of the intelligence available in cable network elements such as the CMTS and CM, as well as plant information to determine type, severity, and location of an impairment. The goal is to proactively correct issues before customers are even aware of them. By sifting through huge amounts of data collected from the network, information can be turned into actionable intelligence to increase the reliability of the network and provide a better experience to cable broadband subscribers.

## 6. Current Activities

As operators became aware of the mass isolation policies and the resulting bandwidth impact on the access network, we needed to respond rapidly. In many cases there was sufficient latent capacity in the network to handle the initial load, particularly in the downstream; but upstream challenges required fast decisions. We are accustomed to a long-term planning cycle of projecting node actions. These actions may include combinations of adding channels for capacity, node segmentation, or physical node splits, and they oftentimes demand forecasting of the bandwidth capacity requirements 6-18 months into the future. As a result of sheltering-in-place, we saw our 18-month forecast occur in a few weeks.

We were fortunate with how resilient the DOCSIS protocol has become, but it does not alleviate the reality of needing to increase spectrum for both upstream and downstream as we progress into the future.

This experience implies that each operator needs to make some difficult decisions quickly. Now that the DOCSIS 4.0 specifications are complete and silicon is being developed for both Extended-Spectrum and Full-Duplex DOCSIS technologies, cable operators need to decide which path to follow. Not necessarily today, but in the near future, as decisions need to be made regarding changing the outside plant to support either direction. Some of those changes may require years of small plant modifications to ultimately support the bandwidth demands of the future.

What we learned from COVID-19 can fill volumes; but the key learning seems to be that even a blind mouse can find a piece of cheese at times. We were very fortunate for the brilliant men and women who developed the DOCSIS infrastructure, as the inherent resiliency of the protocol and products, coupled with the work done in the cable outside plant itself, enabled us to weather the initial storm of a pandemic. However, it is not the time to rest; we need to prepare and prepare rapidly for the next potential event on the horizon.

## 7. Longer-Term Activities

The negative experiences suffered by some MSOs during the COVID-19 bandwidth surge is a harbinger of things to come in the future if MSOs do not continually upgrade their networks to accommodate the ever-growing traffic demands. While Downstream  $T_{avg}$  CAGRs may have slowed from 50% to ~26% in recent years, there is still yearly bandwidth growth occurring in both the Downstream and Upstream directions. In addition, there has clearly been some un-planned bandwidth growth resulting from COVID-19 stay-at-home actions.

Competitive threats from 5G and PON providers may also someday force bandwidth wars, and that struggle will undoubtedly lead to a need for much higher bandwidth offerings within subscriber Service Level Agreements. There will likely come a day when the  $T_{max}$  of 1 Gbps within the Service Level Agreement of today is replaced by a 2 Gbps (or higher)  $T_{max}$  Service Level Agreement in the future. As we have learned in the past, a higher  $T_{max}$  means a higher  $T_{avg}$ . And increases in either or both of those parameters will produce a need for higher Bandwidth Capacities.

To provide adequate Bandwidth Capacity to support these higher  $T_{avg}$  and  $T_{max}$  bandwidths, MSOs are already looking at new technologies for the future. These technologies include the tried-and-true action of running with a Fiber Deeper architecture. Whenever permissible, MSOs will likely pull fiber closer to the home as the years pass to reduce customer counts per node. This effectively reduces the  $N_{sub}$  value within Service Groups, and it also moves the MSOs closer to their end-game technology, which is possibly a Fiber-To-The-Premise or wireless delivery solution.

As a step towards the 10G Initiative of the future, MSOs are also likely to begin deploying the newly defined DOCSIS 4.0 technologies by the middle of the 2020 decade. This may include Full-Duplex DOCSIS (FDX) technologies (with 1.2 GHz Downstreams and up to 684 MHz Upstreams) or Extended Spectrum DOCSIS (ESD) technologies (with 1.8 GHz Downstreams and up to 684 MHz Upstreams). It is quite possible that many MSOs will find a need to utilize only 300-492 MHz of that available Upstream spectrum (which offers ~2.1-3.6 Gbps of Upstream bandwidth capacity).

MSOs are likely to deploy these DOCSIS 4.0 technologies even before they enable them, because they typically want to deploy gear that will last deep into the future. For many MSOs, DOCSIS 4.0 technologies are planned for use well into the 2030's, so the desire to have the equipment in the field for 10-20 years requires that they begin planning for DOCSIS 4.0 equipment deployments as soon as possible. This also means preparing the outside plant by deploying equipment that can support the spectrum required for the same time period. As an example, if an MSO expects that 10 Gbps SLAs

requiring DOCSIS 4.0 technology are required to be enabled by 2032, and if it requires eight years of diligent plant upgrades to permit ubiquitous enablement of DOCSIS 4.0 technology, then DOCSIS 4.0 field deployments in the outside plant would clearly have to begin by 2024.

Since many of these network technologies will be deployed in DAA environments (both Remote PHY and Remote MACPHY), a critical element of those networks is the Converged Interconnection Network (CIN) that provides Ethernet connectivity from the head-end to the node. Today, CIN networks tend to be limited to 10 Gbps Ethernet optics. A Fiber Node supporting two 10 Gbps Service Groups would obviously require more than 10 Gbps of Ethernet capacity. Thus, as bandwidth capacities of DOCSIS 4.0 deployments rise in the future, it is probable that there will be a move to 25 or 50 Gbps Ethernet Optics within the next decade.

If daisy-chaining of nodes becomes popular or if future architectures employ an Ethernet hub near the nodes, then MSOs may also find value in using the newly defined Coherent Optic technologies. These technologies will provide very high bandwidth capacities (>500 Gbps) at a very low cost for the relatively short-hop fiber runs between head-ends and hubs/nodes.

## 8. Conclusion

There were some great learnings that to date have come out of COVID-19 shelter-in-place orders:

1. You can never have too much fiber or spectrum in your network
2. Leverage existing technologies as long as you can, but not longer
3. Not every node is created equal
4. There is an established order to things, but sometimes it can get mixed up

And most importantly...

5. Serendipity and hope are not a business plan

We were very fortunate as an industry that all of the hard work that went into DOCSIS specifications, silicon, and products, over the past 24 years has shown that the capabilities of the DOCSIS ecosystem are what enabled the cable broadband industry to lead the way providing incredible service during a pandemic. As leaders of the cable industry it is incumbent on us to prepare the way for those that will follow in our footsteps, as we have followed in the footsteps of those that came before us. We must therefore create a series of technologies that will allow cable to continue exceeding the demands of our customers for the next 50+ years.

This implies planning, committing to, and supporting those who are working on and developing these new technologies. Keeping our focus on what our customers need today and in the future will lead each of us to make the right decisions for our companies; this means not getting led down a path that may look shiny today but will tarnish rapidly as the reality of long-tail and support costs kick in.

Now is the time to analyze where our companies and customers are heading, plan for our network transformations, develop resources towards executing our plans, and commit to our vendors for our directional choices. Like anything else in cable, these things take time and now is that time to begin.

In conclusion, we need to **act now**, **be bold**, and **stay true** to our individual directions.

## Abbreviations

|          |                                                  |
|----------|--------------------------------------------------|
| ACK      | acknowledgement                                  |
| BIC      | binary increase congestion                       |
| bps      | bits per second                                  |
| CAGR     | compound annual growth rate                      |
| CIN      | converged interconnect network                   |
| CM       | cable modem                                      |
| CMTS     | cable modem termination system                   |
| COVID-19 | Coronavirus disease 2019                         |
| CUBIC    | an enhanced version of BIC                       |
| DAA      | distributed access architecture                  |
| DOCSIS   | data over cable service interface specifications |
| DSL      | digital subscriber line                          |
| HFC      | hybrid fiber-coax                                |
| Hz       | hertz                                            |
| IP       | Internet Protocol                                |
| Gbps     | gigabits per second                              |
| GPON     | gigabit passive optical network                  |
| IPTV     | Internet Protocol television                     |
| ISBE     | International Society of Broadband Experts       |
| Mbps     | megabits per second                              |
| MHz      | megahertz                                        |
| MSO      | multiple system operator                         |
| OFDMA    | orthogonal frequency division multiple access    |
| OTT      | over the top                                     |
| PNM      | proactive network maintenance                    |
| PON      | passive optical network                          |
| SC-QAM   | single carrier quadrature amplitude modulation   |
| SCTE     | Society of Cable Telecommunications Engineers    |
| TCP      | transmission control protocol                    |
| TV       | television                                       |
| QAM      | quadrature amplitude modulation                  |
| QoE      | quality of experience                            |
| QUIC     | a general purpose transport layer protocol       |

## Bibliography & References

1. Traffic Engineering in a Fiber Deep Gigabit World, John Ulm and Tom Cloonan, 2017 SCTE Cable-Tec Expo.

# **A Flexible and Scalable Architecture for Over-the-Air Credentials Provisioning**

A Technical Paper prepared for SCTE•ISBE by

**Alexander Medvinsky**

Engineering Fellow

CommScope

6450 Sequence Dr., San Diego CA 92121

+1 858-404-2367

sasha.medvinsky@commscope.com

**Dr. Tat Chan**

Distinguished System Engineer

CommScope

6450 Sequence Dr., San Diego CA 92121

+1 858-404-3252

tat.chan@commscope.com

**Dr. Xin Qiu**

Senior Director of Engineering

CommScope

6450 Sequence Dr., San Diego CA 92121

+1 858-404-4212

xin.qiu@commscope.com

**Jason Pasion**

Senior Manager of Software Engineering

CommScope

6450 Sequence Dr., San Diego CA 92121

+1 858-404-2241

jason.pasion@commscope.com

# Table of Contents

| <b>Title</b>                                                                                               | <b>Page Number</b> |
|------------------------------------------------------------------------------------------------------------|--------------------|
| 1. Introduction.....                                                                                       | 3                  |
| 1.1. OPUS System Overview .....                                                                            | 4                  |
| 2. Existing OPUS Use Cases .....                                                                           | 6                  |
| 2.1. Provisioning of Credentials based on Operator and Model Authorization .....                           | 6                  |
| 2.2. Provisioning Credentials Based on Proof of Subscription .....                                         | 11                 |
| 2.3. Provisioning of New Credentials with Offline Matching Existing Device ID .....                        | 13                 |
| 2.4. Provisioning of New Credentials with Offline Matching Existing Device ID for Wireless<br>Devices..... | 16                 |
| 2.5. Provisioning of New Credentials with Online Matching Existing Device ID .....                         | 19                 |
| 3. Scalability and Benchmark Results .....                                                                 | 22                 |
| 4. Future Work.....                                                                                        | 22                 |
| 5. Conclusion.....                                                                                         | 23                 |
| Abbreviations and Definitions.....                                                                         | 24                 |
| Bibliography & References.....                                                                             | 25                 |

## List of Figures

| <b>Title</b>                                                                           | <b>Page Number</b> |
|----------------------------------------------------------------------------------------|--------------------|
| Figure 1 - OPUS System Overview .....                                                  | 4                  |
| Figure 2 – Provisioning of Credentials Based on Operator and Model Authorization ..... | 7                  |
| Figure 3 – Provisioning of Credentials Based on Proof of Subscription .....            | 12                 |
| Figure 4: Provisioning of Credentials Matching Existing Device ID .....                | 14                 |
| Figure 5 – Provisioning of Credentials Online Matching Existing Device ID .....        | 20                 |

# 1. Introduction

During its lifetime, a device may need to be updated for a variety of reasons. New network access mechanisms or new types of applications and services may be introduced. This can mean that new device digital identities will need to be installed in already deployed devices to enable such new use cases. Examples of new digital identities include new DRM or conditional access credentials for new sources of content, IoT device certificates, credentials for a new copy protection interface such as DTCPv1, DTCPv2, HDCP 1.x, HDCP 2.x, etc.

Provisioning of credentials into devices already connected to the Internet and deployed to individual subscriber homes cannot rely on network or perimeter security. Even when devices are in an enterprise network or in a network operator's domain, it is still prudent to deploy "defense in depth" by providing end-to-end authentication and encryption all the way to the target device, in addition to any perimeter security such as firewalls, IP address filtering and port mapping. Each device and credentials provisioning server need a well-secured root of trust, delivery of credentials should be secured end-to-end and protected against a variety of network-based attacks.

A provisioning system handling different types of credentials has to support a variety of authorization models for different network operators and content providers. Different authorization interfaces are utilized to validate that a legitimate authorized device is being provisioned and that it belongs to a legitimate subscriber authorized for new credentials.

A credentials provisioning system may require a high degree of scalability for large populations of subscribers of premium content or for IoT appliances. Millions of devices may need to be provisioned with new credentials in a relatively short period of time. The worst-case scalability scenario is probably when a DRM or conditional access system is compromised, and every subscriber requires a new set of credentials within a short time period.

This paper describes an architecture of the CommScope credentials provisioning system called Online PKI Update System (OPUS) that has evolved over 10+ years in order to handle a variety of operator-specific and DRM-specific requirements with reliability, flexibility and scalability in mind.



## 1.1. OPUS System Overview

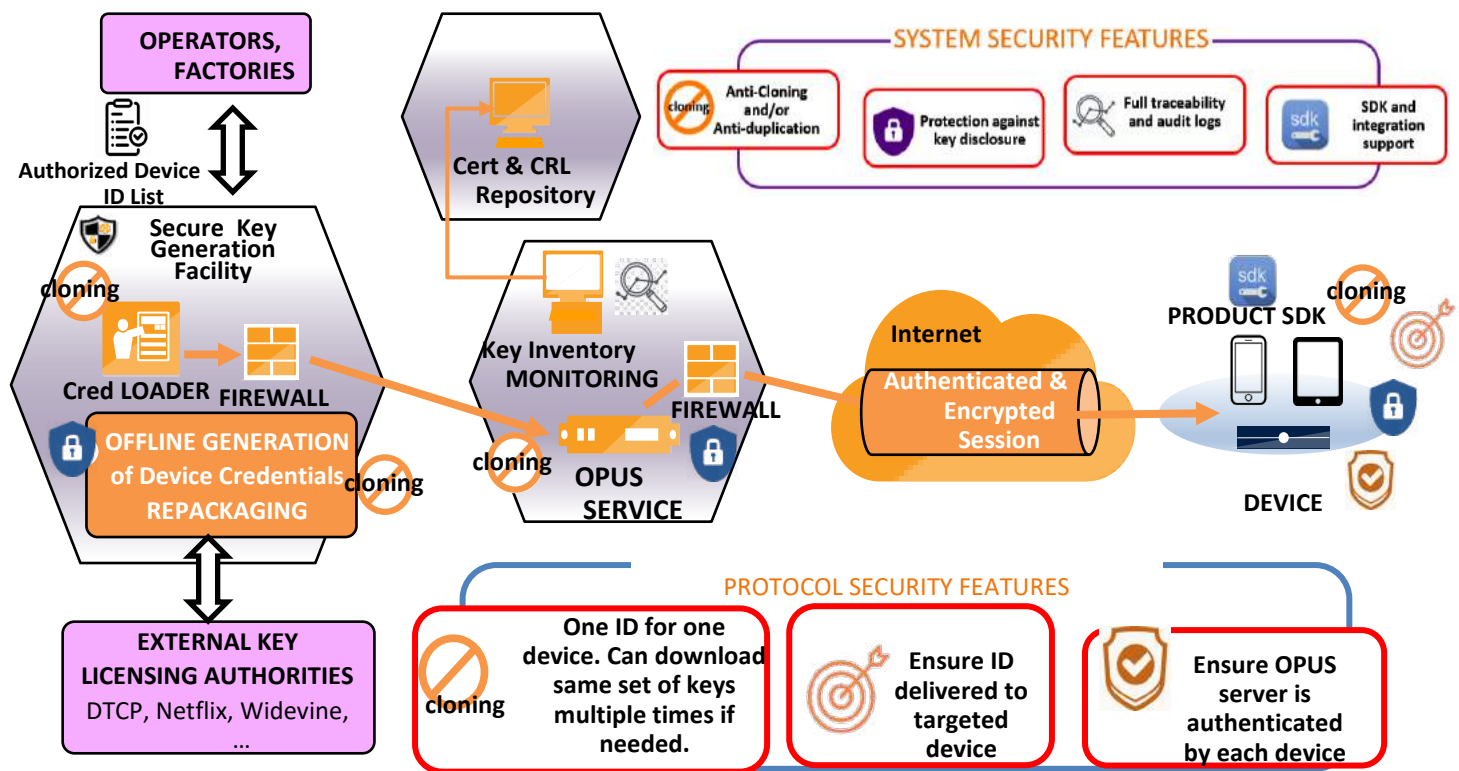


Figure 1 - OPUS System Overview

The OPUS service is accessible by provisioning clients via the Internet. It is protected from malicious attacks by means of a firewall, use of cryptographic specialized hardware, frequent security scanning and patching and other generally recommended security practices. The OPUS service consists of a cluster of frontend machines that are Internet facing to receive and pre-process provisioning requests. Sensitive and secure operations are performed by the backend servers, where the secure database and hardware security modules (HSMs) are located. Additional firewall is installed between the Frontend and Backend to provide additional protection. A hardware-based load balancer is utilized to distribute the workload among the cluster of frontend/backend machine pairs. Device credentials are generated and packaged in an offline secure key generation facility to provide end-to-end security from generation or re-packaging to final consumption on the target device.

There are in general two categories of credentials. In the first category, the credentials are generated locally in the secure key generation facility. It is commonly recommended to generate secret/private keys inside the device itself. However, the quality of a random number generator is highly variable across various device categories and developers that are not skilled in security could utilize a poor random number source such as C functions `rand()` and `srand()`. Even Linux random number source from `/dev/random` and `/dev/urandom` are not considered to be a high-quality random number source and may not be adequate for cryptographic use. Therefore we chose to generate cryptographic material with an HSM inside the secure Key Generation Facility so as to guarantee a sufficient degree of randomness for device credentials. The generated credentials are secured in both the offline Key Generation Facility and

during the transport to the OPUS server and end device so that the risk of compromise prior to the credentials reaching a target device is minimal.

In the second category, the credentials are provided by external parties. The key generation facility in that case only re-packages the credentials to the correct format and applies additional encryption. In some cases, key generation or re-packaging may be based on an authorized device ID list provided by third party such as the Network Operators or Factories. Credentials generated or re-packaged are loaded to the OPUS server through a Credentials Loader application.

The OPUS system utilizes a proprietary request and response message protocol to communicate with provisioning clients. The protocol was evolved over time to provide:

- Flexibility of provisioning any type of device credentials, including but not limited to X.509 device certificates.
- Support for a variety of authorization models some of which that occurred in practice are described in this paper.
- Session security at the application layer with low performance overhead, not requiring additional session setup using a protocol such as TLS or IPsec. Especially considering that each session with an individual device is very short-lived, just for delivery of a single set of device credentials.
- Seamless integration with load balancers, simplified since there is no session setup (OPUS server is stateless)
- Flexibility in terms of adding new cryptographic algorithms

To simplify the integration process, typically an OPUS SDK is built by the CommScope OPUS team and provided to a client device software team. The OPUS SDK handles all the OPUS related processing, including cryptographic operations. This SDK allows fast and easy integration with client teams utilizing the OPUS service and at the same time ensures that the security features on the client side are implemented as designed.

The OPUS solution provides many features crucial to a credential provisioning system, including anti-cloning, strong key protection, reporting, and easy integration. The system is configurable to safely allow the same credential to be downloaded multiple times by the same device in the legitimate cases where the credential on the device may be lost or corrupted. Typically, this is allowed only when that credential can be securely bound to a specific device as was the case in a couple of use cases described in this paper.

In terms of key protection, end-to-end encryption is used in most cases such that device keys are encrypted from generation/re-packaging and decrypted only at the final device. Unique per-device encryption is utilized whenever it is available. In the cases where global encryption is applied to the delivery of device credentials, additional session-based encryption is added by the OPUS server to mitigate the risks of cloning.

Reporting is another important aspect in provisioning. OPUS system maintains logs from key generation/packaging to the actual provisioning transactions, allowing usage reports to be created that may be needed by Network Operators. The logs are crucial for troubleshooting and debugging issues associated with the device credentials provisioning. There is an inventory monitoring service in the system which can generate key inventory reports. When key inventory drops below a preset threshold, a trigger will be created so that responsible parties will be alerted, and more data will be generated or externally acquired and repackaged accordingly.

The OPUS system provides a flexible and scalable solution for secure identity provisioning for devices in the field. In the following, many actual use cases will be discussed in more detail.

## 2. Existing OPUS Use Cases

### 2.1. Provisioning of Credentials based on Operator and Model Authorization

The following is a common use of OPUS for field provisioning of DRM keys into deployed digital set-tops and other secure video rendering devices. Each deployed device has been provisioned with a factory-installed unique credential. Most common type of credential in use is an X.509 device certificate chain, issued by either the network operator or device manufacturer.

Netflix credentials in particular require that an operator signs a business agreement with Netflix and opens a corresponding device account for a specific device model and region. An example of this process is documented here: <https://openconnect.netflix.com/en/#what-is-open-connect>. This process sometimes takes much longer than anticipated and is completed after a particular device model is already in mass production. Therefore, it becomes necessary to provide a secure field provisioning interface to download unique per-device Netflix keys online.<sup>1</sup>

Another common use case is provisioning of Widevine DRM and Attestation Keys into Android and Android TV devices.<sup>2</sup> These Android credentials require for a device vendor to set up a per-model account with Google and when Widevine keys are utilized in conjunction with a Netflix service, then separate accounts for each network operator are also required.

Yet another use case when DRM keys may be provisioned into device following their manufacture is an introduction of a brand-new DRM or copy protection system. DTCP version 2 was introduced in 2017<sup>3</sup> and adds a higher level of content protection which requires a Trusted Execution Environment inside each device and makes use of larger key sizes and stronger cryptographic algorithms than what is defined in DTCP version 1. DTCP version 2 is more suitable for protection of premium 4K and HDR content than DTCP version 1. By the year 2017, there was already a population of secure digital set-top boxes and digital TVs with a sufficient level of HW security and these devices are capable of supporting DTCPv2 and protecting the streaming of 4K and HDR content from a digital set-top box to a digital TV or between server and client set-top boxes. Again, this is a use case for OPUS to update qualified devices with the corresponding DTCPv2 keys.

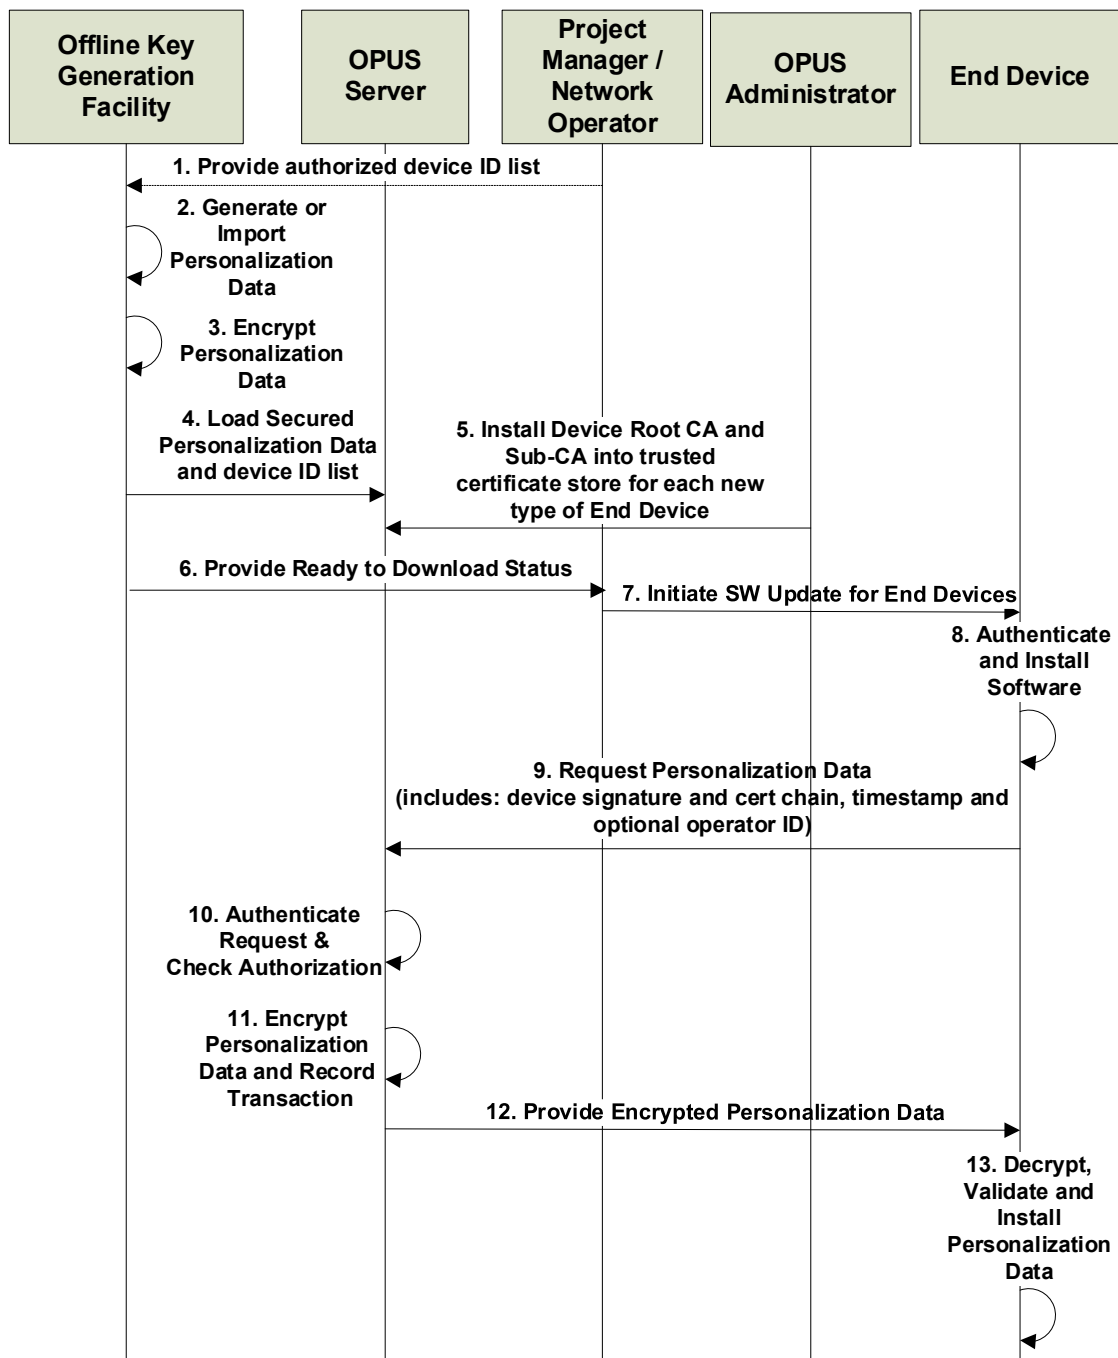
The following sequence diagram demonstrates the use of OPUS for provisioning of new DRM credentials into each device:

---

<sup>1</sup> Linux-based and Windows-based Netflix clients are based on Microsoft PlayReady DRM <https://www.lightreading.com/cable-video/netflix-taps-playready-as-primary-drm/d/d-id/677491> but still require additional Netflix-specific device keys provisioned into each device. PlayReady DRM keys do not require per-operator Netflix accounts and can be provisioned in the factory ahead of time.

<sup>2</sup> See <https://developers.google.com/android-partner/guide/keybox> and <https://developer.android.com/training/articles/security-key-attestation>.

<sup>3</sup> See <https://www.dtcp.com/dtcp.aspx>



**Figure 2 – Provisioning of Credentials Based on Operator and Model Authorization**

The above sequence of credentials provisioning steps are as follows:

1. Network operator optionally provides a list of devices requiring an upgrade and new DRM credentials. If this list is incomplete, then this whole sequence may need to be iterated multiple times.

In majority of cases the operator did not require this list and all devices of a particular model and belonging to a particular operator were allowed to upgrade. But there were some cases when the operator wanted us to upgrade only specific devices based on their list of device identifiers. A device identifier here could be a SOC ID, MAC Address or a Serial Number.

New types of device identifiers may be supported in the future on a case-by-case basis. Most of the time a new type of device identifier can be supported with no software changes to the OPUS system. A device identifier which is a function of a device certificate (e.g., hash or fingerprint of the device certificate) may be introduced in the future as an OPUS enhancement.

2. Generate the new sets of device credentials (e.g., issue X.509 certificates). Or in other cases which we had to commonly support – obtain and then re-package and re-encrypt 3<sup>rd</sup> party DRM credentials e.g., Netflix, Google Widevine or Google Attestation Keys. This may require a SOC-specific format that can be later processed by the chip inside the trusted execution environment (TEE).
3. Encryption may utilize either a global HW key known to the SOC TEE environment or unique per-SOC encryption whenever the SOC ID is known in advance. Decryption is often required to occur inside a TEE on the SOC, especially when the new DRM is expected to handle 4K or HDR high value content.

When the SOC ID is known, it may be utilized to either derive a unique per-SOC key from a global HW key and SOC ID, or utilized to look up a unique One Time Programmable (OTP) key for that SOC in some internal or external database. For CommScope's past utilization of OPUS – we had to handle all of these use cases.

4. Load the generated new DRM credentials onto the OPUS server which is accessible online by devices that are being upgraded. Device credentials were generated in an offline air-gapped facility and so this transfer may be done by a human operator with removable media.
5. At some point prior to a new device model starting to submit requests to OPUS, an OPUS administrator must obtain and install the device root CA and device sub-CA that will be trusted by the OPUS server for submitting requests for device credentials. There can be multiple trusted CA certificate chains configured for the same (operator ID, credentials type) combination since a particular operator may have several different device models with different factory-installed device certificates from different issuers that all require the same new credentials to be downloaded in the field.

This process needs to be repeated for each new credential type – even if it is the same device model requesting a new credential. For more flexible authorization rules, an operator could require the same device model to utilize a different certificate chain for downloading a different type of device credential. There shouldn't be a limit on the number of CAs installed on the OPUS server, covering any number of authorized device models and types of credentials.

Besides configuring the list of trusted CA certificates, there are other miscellaneous configuration parameters that should be set for each credential type. For example, the type of the key agreement algorithm (e.g., Diffie-Hellman or Elliptic Curve Diffie-Hellman), key size or Elliptic Curve ID, etc.

This step is required for all use cases but for simplicity is not shown on subsequent diagrams.

6. A project manager, possibly a member of customer's operations team, is given a go ahead that the new credentials are now ready to be downloaded.
7. The project manager coordinates a device update to prepare devices for a new credentials download. In this example, device update involves an authenticated software update where the new software has the logic to connect to the OPUS server and request the new credentials (in step 8 below). Well ahead of this step CommScope builds an OPUS client SDK for each different device platform to make it easier to complete this software update.

In other cases, devices may already have the latest software capable of connecting to the OPUS server but were waiting for some sort of a trigger such as e.g., a TR-069 configuration message. It is not advisable to wait until the end user tries to access content requiring the new DRM and then begin the credentials download – to avoid any potential delays that are visible to the user.

Section 3 on scalability shows that OPUS is really a cluster of servers capable of very large throughput, but still – there is a lot of unpredictability in the Internet performance and it is best to install DRM credentials in advance of the need to use them.

8. If the software update is required and has not yet taken place, the software is downloaded, verified and installed during this step.

A software update is required under a variety of circumstances such as:

- A credentials update was not anticipated during the initial device manufacturing and the corresponding interface to the OPUS server was not implemented at that time.
- Operator's device authorization mechanisms for new credential updates were not defined at the time of device production or had undergone some changes since then.
- A field update for new credentials was anticipated prior to manufacture, but a strategic decision was made to accelerate device production schedule and defer the interface to download credentials to a future software update.
- Any software in the device that is related to the new credentials incurs a per-device license fee and is therefore deferred to a software update and only for specific end users that subscribe to a specific service.

But in some cases a device may already have the right software to interface to OPUS and this step could be replaced with a trigger message to begin a credentials' download.

Typically, after the software update is validated and installed, the device will reboot, and restart and it may undergo secure boot (not shown on the figure).

9. Device submits a request for new DRM credentials. This request must be authenticated, and we normally rely on factory-installed certificate and private key for that purpose. This device should have implemented secure boot such that all software running in the device is authenticated. It is also recommended to utilize a TEE for this purpose.

Over the years, we had to handle some exceptions where for example a device with no public key credentials was being upgraded with a DRM application that requires a unique device certificate. Sometimes a device has credentials that are by contract restricted to protect only a specific interface (e.g., DTCP-IP or HDCP copy protection or credentials for a specific DRM) and so still – there were no device credentials available to authenticate a request to OPUS.

Under those circumstances, a software update to a device (in steps 6 & 7) included a software-protected signing key which can be mathematically encoded using whitebox techniques and a software protected decryption key which can also be whitebox-encoded. Here, a device signs the request with its software protected signing key if a more secure option is not available. And associated device certificate issued to the corresponding public key is also attached.

This message includes a unique device identifier as a separate parameter and may also be included in the device certificate attached to the request. Device identifier is validated against the authorized device ID list (see the next step).

This request also includes a key agreement public key of the device in order to establish a one-time session encryption key – in addition to encryption that already occurred offline in step 3. The device generates its key agreement public/private keypair and the public key is included in this request message.

Typical key agreement algorithms include Diffie-Hellman (DH) and Elliptic Curve Diffie-Hellman (ECDH). OPUS protocol is configurable and allows for introduction of new key agreement algorithms as needed.

10. Here, the request is authenticated by checking certificate chain and signature of the requesting device. Authorization is also verified using multiple configurable options:
  - Check if Operator ID in the request is authorized for this type of credential
  - Check if a particular device model is authorized for this type of a credential
  - Check if this specific device has its ID in the authorized device ID list
  - Check a CRL to make sure that the factory-installed certificate is not revoked. There may be multiple CRLs, but OPUS will check a specific CRL corresponding to the End Device's certificate in the request. Not all device certificates include a `crlDistributionPoint` extension so the server may need to be configured with the corresponding URL manually.
11. After all authentication and authorization checks passed, the OPUS server will either look up a new device credential based on the authorized device ID, or if allowed by configuration, find the next unbound globally encrypted credential.

The server will generate its own key agreement keypair and include its public key in the response that it is preparing. It will also generate a one-time session key based on server's key agreement private key and the client device's key agreement public key in the request message. This session key is utilized to add an extra layer of encryption, preventing the reuse of the same encrypted credentials in any other OPUS session.

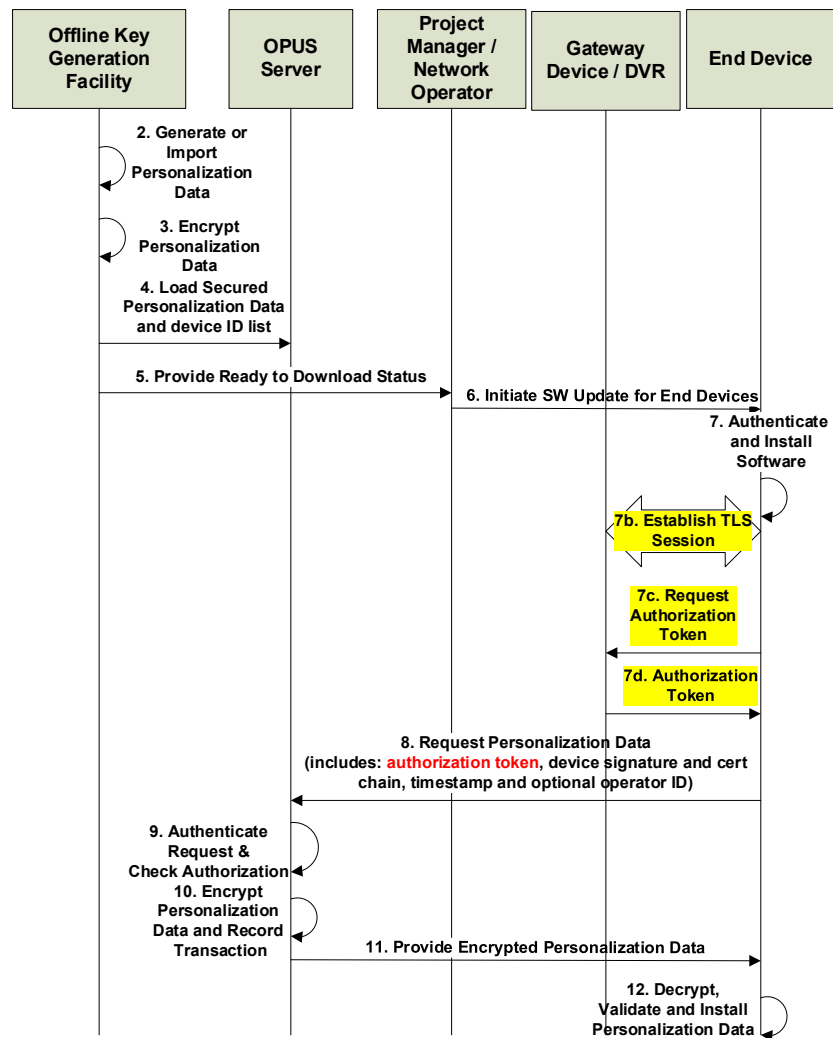
In the case that encryption in step 3 was done with a unique device key, the reuse or copying of those encrypted credentials in another device is already prevented and this extra layer of session key encryption may be skipped for improved performance and scalability.

12. The OPUS server adds its own signature and certificate chain and sends back a reply containing the server's key agreement public key and the encrypted device credentials.
13. Device receives the response, validates OPUS server's signature and certificate chain and then proceeds to decrypt, validate and install its new credentials. These new credentials should be decrypted inside a secure TEE environment and then re-encrypted using a device unique key to bind these new credentials to the specific device.

## **2.2. Provisioning Credentials Based on Proof of Subscription**

Pay TV subscribers may lease or purchase a local DVR that will record subscriber's favorite digital content and this content can later be streamed to subscriber's secondary rendering devices, including additional client set-top boxes, mobile phones and tablets. This subscriber may want to watch content recorded on the main DVR on her tablet connected over a local Wi-Fi network while in a different part of the house from where the DVR is located. Alternatively, a subscriber would like to receive live or pre-recorded content directly to her tablet, cellphone or laptop via a content distribution network while traveling, away from home.





**Figure 3 – Provisioning of Credentials Based on Proof of Subscription**

The above sequence of credentials provisioning steps is similar to the sequence described in section 2.1 with the following differences associated with device authorization:

- Optional step 1 where an authorized device ID list is provided by the operator is not there.
- A subscriber leases a DVR from an operator and pays for an additional service that allows this subscriber to obtain DVR content or even content direct from the operator's content distribution network on some small number of additional personal devices authorized by the operator, including mobile phones, tablets, PCs, etc.
- To prove that this subscriber has a DVR and is authorized for additional content streaming or download services to a personal device, steps 7b, 7c and 7d were added:

- 7b: establish a 2-way authenticated TLS session. The end device and DVR both have pre-installed certificates that can be used for this purpose. DVR would typically be the TLS server and end device is the client.
- 7c: Device requests an authorization token
- 7d: DVR checks subscriber authorization for content sharing with additional devices and if authorized, signs and returns the authorization token. DVR is separately configured by the operator to enable this service.
  - Authorization step in the DVR may include checking the limit on the number of end devices or the DVR may have its own authorized device ID list received from the operator to make sure that this end device is authorized.

Some examples of additional credentials provisioned into the end device include DTCP-IP, X.509 device certificate, Netflix or Widevine. (Netflix and Widevine DRM are typically utilized to secure content delivery direct from the content provider or the network operator but not from the DVR.)

### **2.3. Provisioning of New Credentials with Offline Matching Existing Device ID**

In some cases a network operator may decide to switch to a new Digital Rights Management system that requires new device credentials such as for example device X.509 certificate chain and private keys. Other DRM systems require their own proprietary DRM credentials in a proprietary format that may need to be installed after devices are already manufactured. In other cases, the operator completes their business agreements with a DRM provider after devices already start shipping and so are shipped lacking the necessary DRM credentials.

An operator may require that such a DRM upgrade is done for a specific set of devices and device IDs that are registered in their network and provides the manufacturer with an authorized device ID list requiring such an upgrade. A device identifier attached to the new device credentials may be for example a MAC Address and it may be included as an X.509 certificate subject name attribute as was the case in the CommScope's experience. Another example of provided device identifiers is a list of SOC identifiers which may be used to encrypt each new DRM credential with a SOC-specific symmetric key.

One way to address these use cases is to process an authorized device ID list in a secure offline facility where we have access to all the factory-provisioned device certificates. For each device ID in the authorized list we look up the original factory-provisioned device certificate and use the included public key to encrypt the newly generated DRM credentials of that device. The X.509 certificate that is part of the new credentials is signed in an offline facility using a Certificate Authority approved by the DRM provider and the device ID in the authorized list is included in the new certificate. Encryption is unique to a specific device and is end-to-end, providing a very high level of security. Clear device credentials are not exposed anywhere outside of the secure offline facility and the device itself.

Alternatively, a SOC ID may be used to look up or derive a SOC-specific and hardware-protected symmetric key and utilized to encrypt the newly generated device credentials. In this case, encryption will still be unique to a specific device and is end-to-end, providing the same high level of security.

CommScope has utilized the deployed OPUS system to perform such device upgrades for multiple (4+) operators which made a decision to switch their whole network from Digital Rights Management System #1 to Digital Rights Management System #2 or in other cases concluded their DRM agreements after devices were already manufactured and delivered to that operator. In each case, the DRM upgrade was successfully completed according to plan.

The following diagram focuses on a use case where each new credential is an X.509 certificate and private keys are encrypted using a public key from a factory-installed certificate:

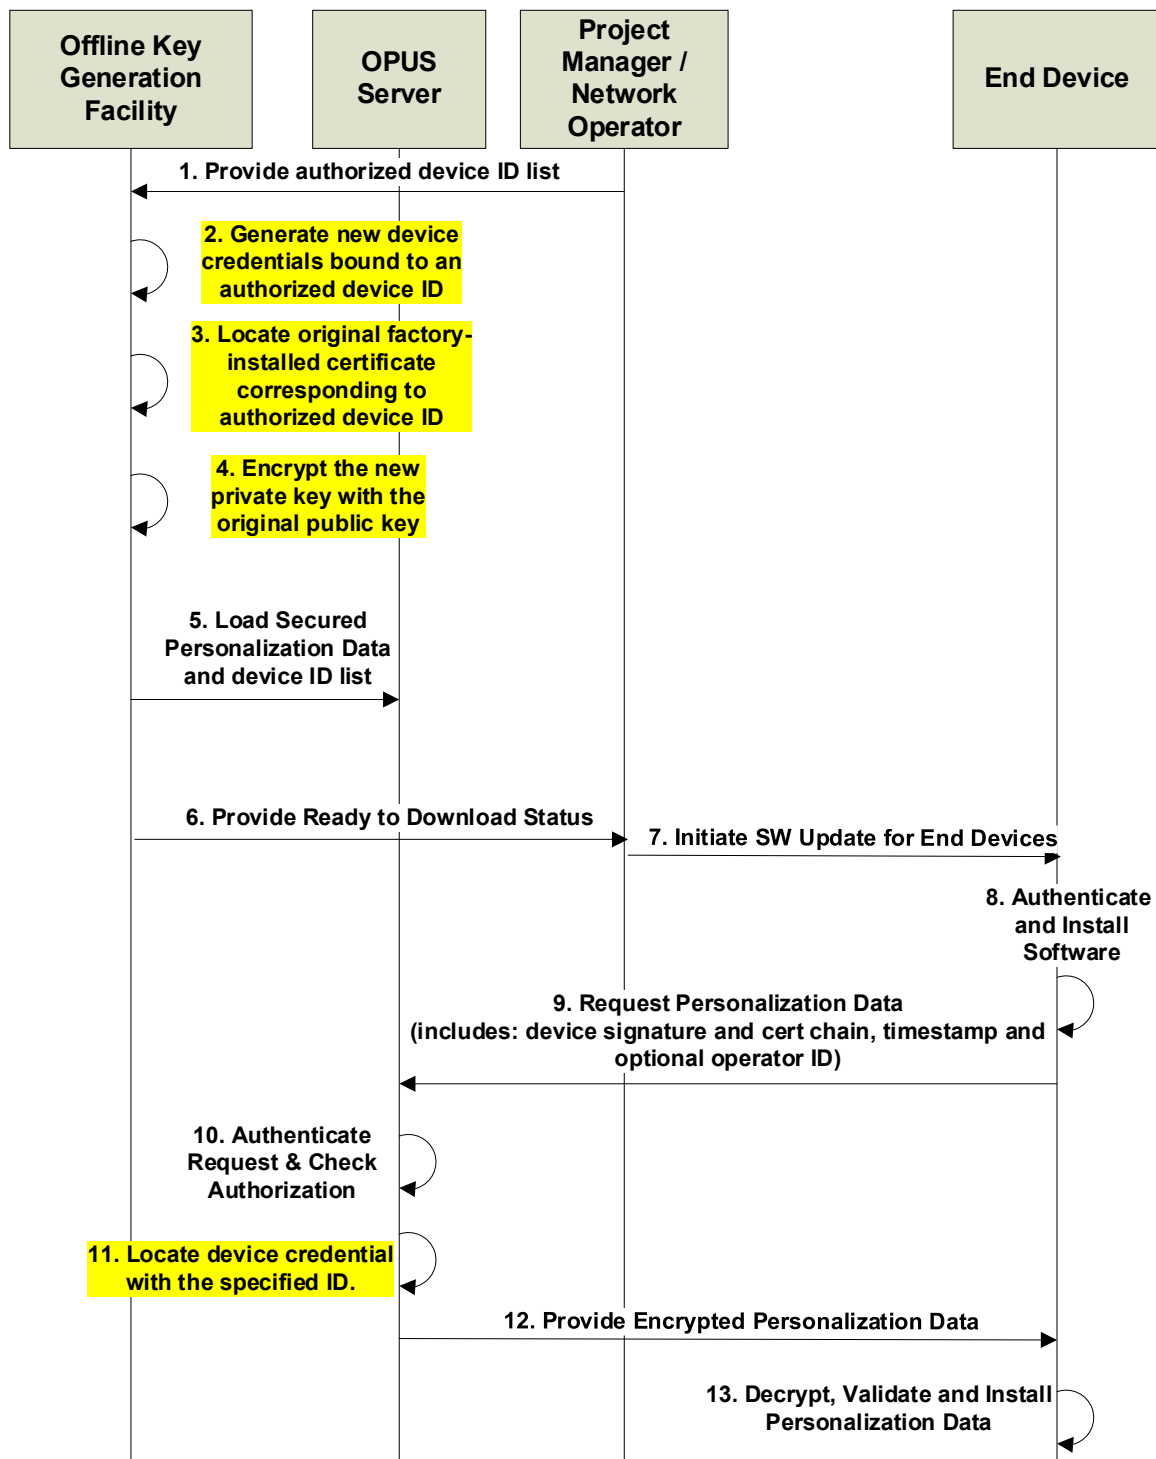


Figure 4: Provisioning of Credentials Matching Existing Device ID

1. Network operator provides a list of authorized device identifiers such as MAC Addresses that are required to be present in the new device certificates.
2. Generate the new device credentials consisting of a public/private keypair and a device certificate chain. The newly generated device certificate includes one of the authorized device identifiers on the list. Certificate is signed using a Certificate Authority that is authorized for use with the new DRM system.
3. Locate the original factory-installed certificate in an offline database. The original certificate may contain the same device identifier. If the original certificate contains a different identifier, then the authorized device ID list provided in step #1 needs to include pairs of (factory-cert-device-ID, new-cert-device-ID) for each device. That way, factory-cert-device-ID can be used to locate the original factory-installed device certificate. Private key corresponding to that factory-installed certificate is not needed.
4. A public key is extracted from the factory-installed certificate and utilized to encrypt the new private key, generated as part of the new device credentials for the new DRM. If RSA public keys are utilized, due to size limitations the RSA public key is utilized to indirectly “wrap” the larger RSA private key structure. A standard RSA wrapping mechanism is specified in [5].
5. Uniquely encrypted device credentials for each ID in the authorized device ID list are loaded to the OPUS server. Encryption was performed in an offline facility and therefore this step requires manual action such as transfer of the keys on removable media. Authorized device ID can be implicit – each authorized device ID is included in one of the uploaded device credentials.
6. A project manager, possibly a member of customer’s operations team, is given a go ahead that the new credentials are now ready to be downloaded.
7. The project manager coordinates a device update to prepare devices for a new credentials download. This may be done via an authenticated software update to the end device. (See a more detailed description on step #6 in section 2.1.)
8. If device received a software update needed for the DRM credentials update in step #7, then the software update is authenticated in this step prior to installation.
9. The device submits a request for new credentials to the OPUS server. In this use case, the device includes the new device ID that it requires in the new credential and the request is signed using the factory-installed device certificate and private key.
10. The OPUS server verifies the signature and certificate chain of the request. Some authorization may be performed here, such as verifying the network operator’s ID in the request to make sure the operator is authorized for this DRM update.
11. Locate the new device credentials matching the device ID in the request. This provides additional implicit authorization – devices that did not have their device ID included in the list provided during step #1 will not have their certificate loaded on the OPUS server. Either that means that device is not authorized for the DRM upgrade – or the operator simply missed providing the identifier of this device.

12. The encrypted device credentials located on the OPUS server are returned to the device in this step. It isn't necessary to apply additional session-based encryption here since the new credentials are encrypted for a specific target device. But the message is signed and anti-replay protection may be needed in the absence of a key agreement protocol such as DH or ECDH.
13. Device receives, verifies and installs its new credentials. Decryption and re-encryption are likely required in order to store the new DRM credentials using DRM-specific encryption. But in some cases it may be sufficient to save the encrypted private key in the exact format as it was received.

In practice, it turned out that it was very hard for operators provide a complete authorized device ID list in step #1 for all devices that required the DRM upgrade. Some devices may have their device identifier modified in repair. Or the operator may be missing some identifiers of newly shipped devices for which they have not yet received a shipping notice. In our practice, this is an iterative process and for each operator we received the authorized device ID list in 5-10 parts and so this whole process was repeated 5-10 times per operator.

One way to handle this upgrade is to have the OPUS server keep track of errors of requested device IDs that were not found in its list, deliver this list of missing IDs to the operator which would then determine if those are legitimate devices that were missed during the previous DRM upgrade iterations. And then repeat the whole process with the previously missing device IDs.

OPUS system has also been utilized in very similar use cases as the above with a few variations, where:

- The new DRM credentials are in a DRM-proprietary format and may not include X.509 device certificates
- Authorized device ID list consists of SOC IDs and DRM credentials are encrypted with a SOC-specific symmetric key instead of a public key from a factory-installed certificate

The sequence of entity interactions looks the same as in the above diagram with minor variations in the type of device identifiers and encryption that was deployed.

## **2.4. Provisioning of New Credentials with Offline Matching Existing Device ID for Wireless Devices**

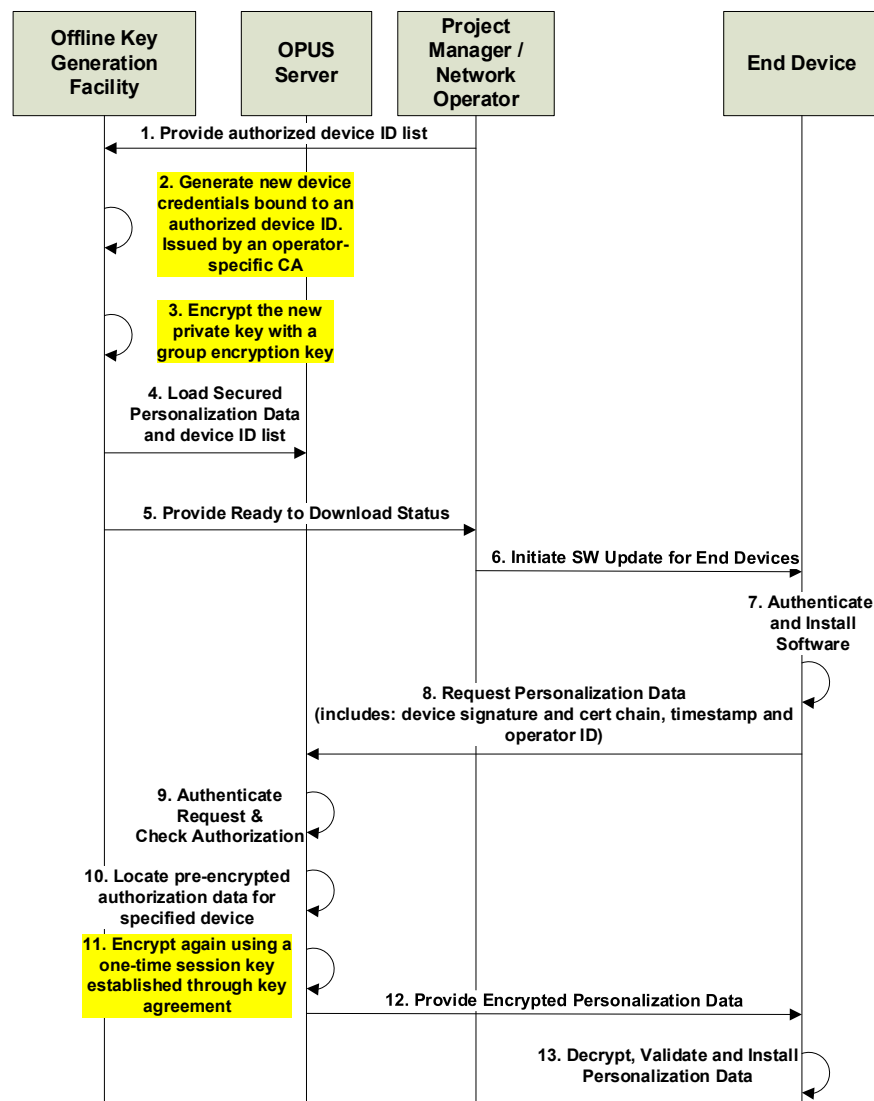
In some cases, a wireless network device such as a router, base station, radio point or radio point controller requires an X.509 device certificate that is utilized in a lower layer protocol security such as for Thread, IKE/IKEv2/IPsec or MACsec. In those cases, it is common to require a device identifier such as the MAC Address to be present in those certificates. Such certificates can be installed in the factory, except when an operator wants to install their own credentials that chain to an operator-specific Root CA.

A device manufacturer may not know network operator's identity prior to manufacture or equipment may change ownership between different operators when a portion of operator's network is sold to a new operator. Such use cases can also be addressed with OPUS providing field provisioning of new device credentials in a very similar manner as the use case described in section 2.3, but with some simplifications.

Encryption in the offline facility is done with a group key which doesn't require locating the original factory-installed certificate. Instead, the OPUS server performs a key agreement protocol to generate a unique session key for an extra layer of encryption thus preventing cloning attacks and reuse of the same encrypted credentials in multiple devices.

The use case below was in practice considered for a wireless network device upgrade, although the same OPUS profile can be applied to field provisioning of any kind of new credentials – when each new credential is targeted to a specific device and a list of the corresponding device identifiers is known in advance. It has some similarity to the use case describe in section 2.3, except that here the OPUS operator may not have access to the factory-provisioned device credentials.

This is illustrated in the following diagram:



The above sequence of credentials provisioning steps is described in more details below:

1. Project Manager or Network Operator provides an authorized device ID list (list of device identifiers for devices requiring new credentials). Each such list is associated with a specific Network Operator, used to determine which CA is used to issue the corresponding device certificates. In practice, this step is iterated periodically. For example, the project manager sends a list for all new devices shipped within a pre-determined time period.

2. Generate the new sets of device credentials (e.g. X.509 certificates) with the device identifiers provided in the list. The CA used to issue the certificates may be specific to the Network Operator.
3. Encrypt each of the device private keys using a group/model-specific symmetric key or a key derived from such a key.
4. Load the authorized device ID list and new credentials onto the OPUS server. The new credentials are identified by a Personalization Type ID that is specific to the type of credential and Network Operator combination. Since the credentials were generated in an offline air-gapped facility, data loading may be performed by a human operator using removable media.
5. The Project Manager or Network Operator is notified that the credentials have been loaded and therefore ready to be consumed by the devices on the list.
6. This step describes the process of the Project Manager or Network Operator triggering the software update to the device for requesting the new credential. This step is the same or similar to step 6 in the section 2.1.
7. In this step, the device verifies and installs the software update necessary to interface to the OPUS server, if it has not been done yet.
8. Similar to the previous use cases, the device submits a request to the OPUS server for new credential, identified by a Personalization Data Type ID. Since the new private key to be provisioned is globally encrypted, a unique session key will be created by the server (in a later step) to mitigate the risk of cloning. Therefore the device generates a key agreement key pair for the request and includes the key agreement public key in the request. In addition, the device also includes an Operator ID for additional validation by the server. The request includes a device signature and a factory-installed certificate.
9. Upon receiving the provisioning request from a device, the OPUS server verifies the signature on the device using the attached device certificate extracted from the request. The server makes sure that the device certificate used is chained to a pre-configured CA for that Personalization Data Type ID. Next, the server extracts the Device ID from the device certificate and verifies if it is on the authorized device ID list. The server would reject the request if any of the verification steps fail.
10. Next, the OPUS server looks up the credential pre-generated for that particular device, based on the Device ID. The credential includes the new device certificate and globally encrypted private key.
11. Since the private key is encrypted globally, OPUS server applies another layer of encryption on top of it using a session key derived from the key agreement algorithm. To do this, the OPUS server extracts the key agreement public key from the request, and generates its own key agreement key pair.
12. The OPUS server sends a response to the device including the new credential (the new device certificate and corresponding private key, which is now double encrypted with the session key), and the server's key agreement public key, needed for the device to derive the same one-time session key. The server also signs the response using its server certificate.

13. After receiving the response, the device verifies the signature of the message using the provided server certificate. The device also makes sure the server certificate chains to a trusted CA embedded in the provisioning software installed earlier. The device then derives the one-time session key using the provided server key agreement public key and its own private key. The device then removes both layers of encryption over the private key, utilizing the one-time session key and then the global encryption key. At this point the device may perform further validation to make sure the certificate is properly formatted and chained to the expected Root CA certificate. It also makes sure the private key and certificate match each other. After all the validation, the device can proceed to store the new credentials.

To avoid cloning, the private key must be stored encrypted using a unique key tied to the device rather than using a global key. The device derives such a unique key from a global key and a device or SOC identifier. Or a unique device key may have already been provisioned (e.g. into secure OTP memory). The device uniquely encrypts the new private key and saves it into persistent storage along with the associated device certificate.

This use case has the security advantage that the signing CA is hosted in the offline key generation facility and not on OPUS. However, a drawback is that system response time may not satisfy the operational requirements imposed by the Network Operator. After the authorized device ID list is provided to the offline key generation facility, there will be a lead time for the new credentials to be available for a device to download. In some cases, Network Operator may impose a short turnaround time, say, one day, which may not be satisfied easily with the many steps (some manual) for this process. The next use case can be used in such scenarios.

## **2.5. Provisioning of New Credentials with Online Matching Existing Device ID**

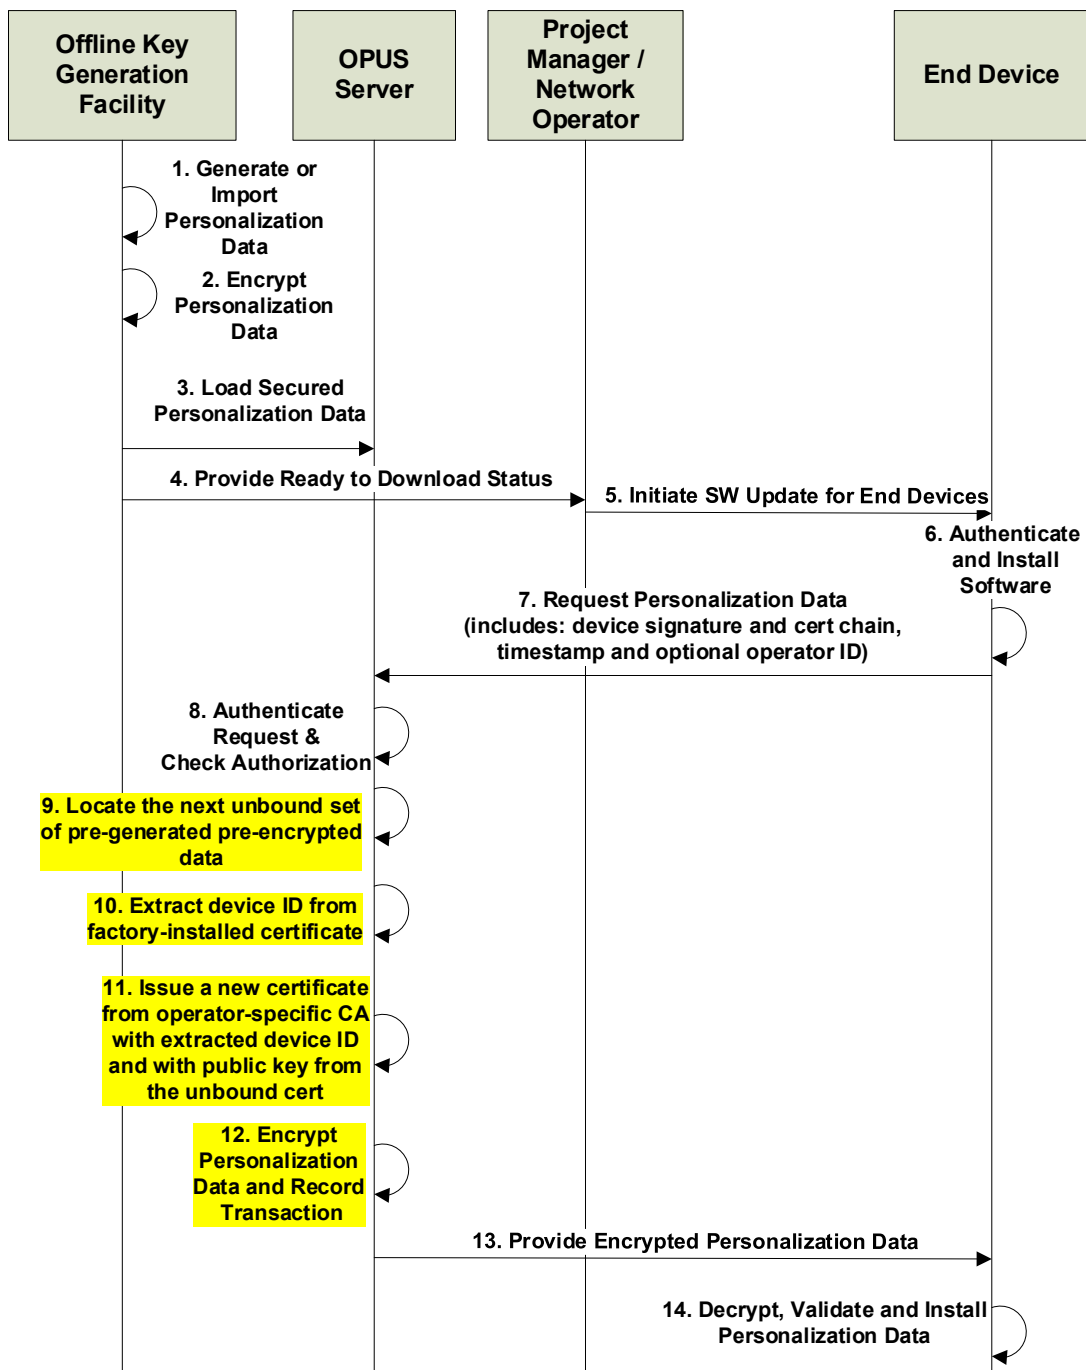
The approach for wireless device credentials provisioning that is described in section 2.4 was seriously considered but as it turned out – after an operator submits a new order for wireless devices, the new certificates have to be made available almost immediately. While it is a more secure approach in section 2.4 where none of the Certificate Authorities are operated online, it did not meet the practical performance constraints.

Therefore, in practice we addressed this use case with a different system architecture that enables the OPUS server to issue new operator-specific certificates. The OPUS server is sufficiently hardened where the certificate issuer's private key is protected in the HSM at all times and device private keys do not need to be exposed on the server. The OPUS server is also sufficiently hardened and is protected behind a firewall.

Just as in the previous section, this use case is not limited only to the wireless device credentials, although in practice that's where it was applied.

This use case is illustrated in the diagram below:





**Figure 5 – Provisioning of Credentials Online Matching Existing Device ID**

The above sequence of credentials provisioning steps is explained in the following:

1. A set of device credentials, called the Origin Credentials, are pre-generated in the offline key generation facility. Each record of Origin Credential includes a device private key and a corresponding Origin Certificate issued using an identifier that is not the final device identifier.

This is because the list of authorized device IDs is not yet known at this point. The generation may be triggered by production volume estimation provided by a Project Manager.

2. The device private key is encrypted just as in Section 2.4, using a group/model-specific symmetric key or a key derived from such a key. This allows the device private keys to be pre-encrypted end-to-end from generation to the device. The OPUS server does not need to have access to the clear device private key and therefore the exposure of the device private keys is minimized.
3. The Origin Credentials are loaded to the OPUS server. These credentials are identified by a Personalization Type ID that is specific to the type of credential but not necessarily specific to the Network Operator. Note that at this point, the credentials are not yet bound to a specific device.
4. The Project Manager or Network Operator is notified that the credentials have been loaded, and provisioning can proceed.
5. This step describes the process of the Project Manager or Network Operator triggering the software update to the device for requesting the new credential. This step is the same or similar to step 6 in the section 2.1.
6. In this step, the device verifies and installs the software update necessary to interface to the OPUS server, if it has not been done yet.
7. In this step, the device submits a request to the OPUS server similar to step 8 of section 2.4. Again, since the device private key is globally encrypted, a one-time session key will be used to add an additional layer of encryption. Note also that in this use case, the request must include the Operator ID, which will be used by the server to determine which CA to use for issuing the final certificate.
8. Upon receiving the request, the OPUS server authenticates the request and checks for authorization, similar to the previous case. However, as an authorized device ID list is not utilized here, the OPUS server only verifies that the factory-provisioned device certificate in the request is issued by a trusted CA pre-configured for this specific Personalization Data Type and for the specified Operator ID.
9. OPUS server retrieves the next available unbound set of Origin Credentials, which includes an Origin Certificate and a corresponding globally-encrypted device private key.
10. OPUS server then extracts the device ID from the factory-installed certificate in the request. This device ID will be used in the final certificate.
11. OPUS server issues a final certificate for the device. The final certificate has the extracted device ID in the subject name, and the public key extracted from the Origin Certificate. The final certificate is signed by the OPUS server using a CA determined by the Operator ID provided in the request. This final certificate is now bound to the device. For additional security, the CA private keys are protected in an HSM and therefore never exposed in the clear on the server. Additionally, an operator-specific certificate template may be selected based on an Operator ID. For example, certificate validity period and some subjectName attributes may be determined by an operator-specific template.

12. In this step, the server derives a one-time session key to encrypt the globally-encrypted device private key. This is the same as in the previous use case. The session key can be derived inside the HSM as well to minimize its exposure.
13. The final credential is then sent to the device.
14. The device decrypts, validates, and installs the credentials the same way as in the previous use case.

In practice, this is the selected method utilized by OPUS to provision operator-specific credentials into wireless devices. While still maintaining end-to-end encryption of private keys from the offline key generation facility all the way to the target device, this method enables real-time provisioning of the device credentials and doesn't depend on manual pre-processing and installation of an authorized device ID list.

### 3. Scalability and Benchmark Results

With a cluster of 2 front end/back end OPUS server pairs, when key agreement is not required (when unique-per device encryption is applied end-to-end), the observed throughput is 30 million transactions per day. The load balancer allows OPUS to increase scalability if necessary and multiple OPUS clusters with their own load balancers may be deployed for different operators or different types of device credentials. The observed performance included hardware acceleration for RSA signatures provided by a hardware security module (in addition to the signing private key protection)

Key agreement protocols such as DH or ECDH may require additional OPUS server pairs to maintain the same level of scalability. This is highly variable depending on the computing power, hardware-based crypto acceleration and a choice of the key agreement algorithm. According to [6], software implementations for ECDH utilizing NIST P256 curve are 3 times faster than 2048-bit classic Diffie-Hellman. ECDH utilizing Curve25519 gains a performance improvement of another factor of 3 over the NIST P256 curve.

For the OPUS design, security of device credentials is the highest priority since additional layers of encryption or larger key sizes can always be offset by additional pairs of FE/BE servers which can be shared across multiple types of device credentials and operators.

### 4. Future Work

The use case described in section 2.5 provides a lot of flexibility to network operators to provision a device with their own credentials that utilize an operator-specific CA and an operator-specific certificate template. However, that use case requires a device to store a new private key that is downloaded into the device from the OPUS server. And that may not be possible for devices that utilize secure OTP memory to store a private key which cannot be modified following manufacture.

For such future use cases, the OPUS server may issue a new operator-specific certificate for the same public key that is part of a certificate installed at manufacture time and corresponds to a device private key that is burned into fuses and is not updateable. Allowing multiple simultaneously valid certificates for the same public key is not a recommended security practice and in this case we should revoke the previously issued certificate right before issuing a new operator-specific certificate.

Another way to address the limited storage for private keys in the OTP security fuses would be to derive a public/private keypair during device startup from a seed value that is locked in OTP. This can be easily done with the Elliptic Curve crypto system where a private key is a random value and is much more difficult with RSA. With this approach, each new operator-specific certificate can be issued to a new public key derived based on the new operator ID and it will not be necessary to revoke a previous certificate unless it was compromised.

Additional OPUS system flexibility will be addressed with future enhancements, including for example:

- Improving algorithm agility through the following means:
  - o OPUS request will include a list of client-supported algorithms and the OPUS server will select the algorithms from that list to protect the corresponding response containing new credentials.
  - o In the case that OPUS server finds client's selection of algorithms inadequate, a client will be automatically redirected to a software update with additional/stronger algorithms. After the client has been upgraded, it will retry.
- Introduction of additional types of device identifiers such as for example a certificate fingerprint

## 5. Conclusion

Today's consumers may own or lease a wide variety of digital entertainment devices all capable of receiving DRM-protected digital content. Different network operators and over the top content services are protected with a variety of DRM systems and a new use case for a particular device model may not be apparent until after those devices are manufactured and delivered to end consumers.

In order to provide maximum entertainment value to consumers, it becomes necessary to provision their devices that are already installed in their home with new DRM credentials. This can be both a daunting scalability challenge and a security concern.

There are many other types of devices such as routers, wireless network access points, radio points and radio point controllers that have similar needs to obtain new cryptographic device credentials after each such device has already been fielded. Such devices may reside in consumer's home network, a private enterprise network or in the network provider's infrastructure. And they may require different types of credentials to be field downloaded.

Furthermore, a credentials provisioning system has to be constantly revised in order to support new use cases involving new authorization models that may be specific to a network operator or to a new industry. This paper demonstrates that a single field provisioning system can be designed with sufficient flexibility and scalability to address such challenges and to handle many use cases.

# Abbreviations and Definitions

|                   |                                                                                                                                                                               |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AES               | Advanced Encryption Standard                                                                                                                                                  |
| BE                | back end (server)                                                                                                                                                             |
| CA                | Certificate Authority                                                                                                                                                         |
| DH                | Diffie-Hellman                                                                                                                                                                |
| DRM               | digital rights management                                                                                                                                                     |
| DTCP              | Digital Transmission Content Protection                                                                                                                                       |
| DVR               | digital video recorder                                                                                                                                                        |
| FE                | front end (server)                                                                                                                                                            |
| ECDH              | elliptic curve Diffie-Hellman                                                                                                                                                 |
| HDCP              | High-Bandwidth Digital Content Protection                                                                                                                                     |
| HDR               | high dynamic range (high quality video content format supported by Ultra HD TVs)                                                                                              |
| HSM               | hardware security module                                                                                                                                                      |
| IETF              | Internet Engineering Task Force                                                                                                                                               |
| IoT               | Internet of things                                                                                                                                                            |
| IPsec             | Internet Protocol Security                                                                                                                                                    |
| ITU-T             | International Telecommunication Union Telecommunication Standardization Sector                                                                                                |
| MAC               | media access control                                                                                                                                                          |
| MACsec            | media access control security (specified by the IEEE 802.1AE-2006 standard)                                                                                                   |
| OPUS              | Online PKI Update Service                                                                                                                                                     |
| Origin Credential | Device credential generated in the factory which requires at least partial transformation and replacement on the OPUS server prior to being provisioned into a target device. |
| OTP               | one time programmable                                                                                                                                                         |
| PKI               | public key infrastructure                                                                                                                                                     |
| RFC               | request for comments (an IETF standards document)                                                                                                                             |
| RSA               | Rivest-Shamir-Adleman (public key cryptosystem)                                                                                                                               |
| SDK               | software development kit                                                                                                                                                      |
| SOC               | system on a chip                                                                                                                                                              |
| SSL               | Secure Sockets Layer                                                                                                                                                          |
| Sub-CA            | subordinate certificate authority                                                                                                                                             |
| TEE               | trusted execution environment (e.g., ARM TrustZone, or a separate security processor/co-processor with its own memory space)                                                  |
| TLS               | Transport Layer Security (IETF-standardized successor to SSL)                                                                                                                 |
| X.509             | Standardized format for digital public key certificates. It is published as an ITU-T standard and has a corresponding IETF standard, RFC 5280.                                |

## Bibliography & References

| Ref | Document Name                                                                                                                                                                                                               | Authors                                                                   | Version    |
|-----|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------|------------|
| [1] | Netflix Open Connect,<br><a href="https://openconnect.netflix.com/en/#what-is-open-connect">https://openconnect.netflix.com/en/#what-is-open-connect</a>                                                                    | Netflix                                                                   |            |
| [2] | Keybox User's Guide,<br><a href="https://developers.google.com/android-partner/guide/keybox">https://developers.google.com/android-partner/guide/keybox</a>                                                                 | Google                                                                    |            |
| [3] | Verifying Hardware Backed Keypairs with Key Attestation,<br><a href="https://developer.android.com/training/articles/security-key-attestation">https://developer.android.com/training/articles/security-key-attestation</a> | Google                                                                    |            |
| [4] | DTCP and DTCP2 specifications,<br><a href="https://www.dtcp.com/dtcp.aspx">https://www.dtcp.com/dtcp.aspx</a>                                                                                                               | DTLA                                                                      |            |
| [5] | Use of the RSA-KEM Key Transport Algorithm in the Cryptographic Message Syntax (CMS), IETF RFC 5990                                                                                                                         | J. Randall,<br>B. Kaliski,<br>J. Brainard,<br>S. Turner                   | Sept, 2010 |
| [6] | eBACS: ECRYPT Benchmarking of Cryptographic Systems,<br><a href="http://bench.cr.yp.to/call-dh.html">http://bench.cr.yp.to/call-dh.html</a>                                                                                 | VAMPIRE<br>(Virtual<br>Application and<br>Implementation<br>Research Lab) |            |

# **Leveraging Legacy Video In Digital Access Architecture Networks**

A Technical Paper prepared for SCTE•ISBE by

**Wesley Weiss**

Network Architect III, Video Infrastructure  
Shaw Communications Inc.  
2728 Hopewell Place NE, Calgary, Alberta, Canada T1Y 7J7  
587-393-0654  
Wesley.Weiss@sjrb.ca

**Anjan Bajwa**, Co-writer/Shaw Communications Inc.

**Corwin Martens**, Contributor/Shaw Communications Inc.

# Table of Contents

| Title                                                                                       | Page Number |
|---------------------------------------------------------------------------------------------|-------------|
| Table of Contents .....                                                                     | 2           |
| 1. Introduction.....                                                                        | 4           |
| 2. Drivers for Legacy Video Integration into DAA .....                                      | 4           |
| 3. Principal and Auxiliary Core Overview .....                                              | 5           |
| 4. Current State of Legacy Video (Pre-DAA) .....                                            | 5           |
| 5. DAA Technology and Components.....                                                       | 7           |
| 5.1. VUE Solution.....                                                                      | 8           |
| 5.2. APEX 3000-NE.....                                                                      | 10          |
| 5.3. Video Topology Manager .....                                                           | 10          |
| 5.4. Legacy SCTE55-1 Interactive Signaling .....                                            | 10          |
| 5.5. DSG SCTE55-1 Interactive Signaling .....                                               | 13          |
| 5.6. Converged Interconnect Network.....                                                    | 13          |
| 5.7. Timing.....                                                                            | 14          |
| 5.8. Remote PHY Device .....                                                                | 14          |
| 6. DAA Monitoring Solution .....                                                            | 15          |
| 7. Video and VOD Encryption .....                                                           | 15          |
| 8. Video Broadcast Acquisition Architecture .....                                           | 16          |
| 9. Narrowcast Architecture .....                                                            | 17          |
| 10. IPv6 Addressing Standards.....                                                          | 19          |
| 11. Multicast, DEPI, and SessionID Relationships with Physical Frequencies on the RPD ..... | 19          |
| 12. Regional Video Realities and Solutions .....                                            | 20          |
| 13. Shaw Automation Strategy.....                                                           | 20          |
| 13.1. Service Director.....                                                                 | 21          |
| 13.2. Resource Inventory .....                                                              | 22          |
| 13.3. VOD Central Database .....                                                            | 22          |
| 13.4. Automating the Video Topology Manager.....                                            | 23          |
| 14. RPD Provisioning and Activation Process Through Automation .....                        | 23          |
| 15. Lessons Learned.....                                                                    | 25          |
| 16. Conclusion.....                                                                         | 26          |
| Terms and Abbreviations .....                                                               | 27          |
| References.....                                                                             | 30          |



## List of Figures

| <b>Title</b>                                                                            | <b>Page Number</b> |
|-----------------------------------------------------------------------------------------|--------------------|
| Figure 1 Legacy Hub Site Video Architecture .....                                       | 6                  |
| Figure 2 DAA Data Flow .....                                                            | 7                  |
| Figure 3 Shaw DAA Rack Layout .....                                                     | 9                  |
| Figure 4 RPD and vARPD Relationships .....                                              | 11                 |
| Figure 5 DAA Interactive Network (Physical Locations) .....                             | 12                 |
| Figure 6 DAA Work Flow.....                                                             | 12                 |
| Figure 7 DSG Interactive and Data Flow Traffic .....                                    | 13                 |
| Figure 8 CIN Network.....                                                               | 14                 |
| Figure 9 DAA Monitoring Architecture .....                                              | 15                 |
| Figure 10 Broadcast APEX 3000-NE Data Centre Network Configuration .....                | 16                 |
| Figure 11 NE GUI Showing VOD Encryption.....                                            | 17                 |
| Figure 12 Narrowcast Architecture of the APEX 3000-NE Blade and VUE Configuration ..... | 18                 |
| Figure 13 SD DAA Interactive Map.....                                                   | 22                 |
| Figure 14 Activation Sequence Diagram During the RPD Activation Process .....           | 24                 |

# 1. Introduction

Digital Access Architecture (DAA) is driving the transformation in next-generation cable-access networks. In contrast, the effective integration of legacy video services and automation deployments can be a daunting, and an often-overlooked task. DAA technology is continually evolving, and these issues apply to many operators facing similar challenges. Creating a solid foundation and adopting DAA technology in stages allows for a scalable video architecture and a *one-touch* approach to deployment, thereby avoiding expensive upgrade costs to customers and infrastructure.

This whitepaper explores Shaw's operational and technical complexities of integrating legacy video into a DAA network and presents the best practices and lessons learned from this undertaking. A robust discussion around functional and automation transformations are covered to prepare operators for large-scale Remote PHY Device (RPD) deployment, modification, and monitoring in an evolving network ecosystem.

## 2. Drivers for Legacy Video Integration into DAA

Shaw believes that DAA is the path forward in upgrading to a robust, resilient, high performance, future-proof, hybrid fibre-coaxial (HFC) network. However, a large install base of legacy customer premise equipment (CPE) exists and needs to be considered. Shaw set out to create an architecture that fully supports in-place capabilities of legacy CPE while still being able to embrace a move to DAA, allowing for upgrading to DAA infrastructure without the added cost of CPE replacement. Furthermore, due to the realities of hub site to node combining, a non-trivial amount of power and space is being expended to the end of manual radio frequency (RF) combining. The ability to reduce entire racks of connected gear to a handful of 1RU switches and dense wavelength division multiplexing (DWDM) hardware could offer significant power, space, and cooling savings. As a result, reducing the requirements for real estate, possibly being able to house an entire hub site in a small fibre cabinet.

While initially researching architectures for legacy Video DAA, we endeavoured to utilize a solution that could re-use as much of the original legacy acquisition, encryption, and out-of-band (OOB) communication hardware. Leveraging as much virtual and Commercial Off-the-Shelf (COTS) hardware as possible to minimize vendor lock-in and maximize flexibility in the event of changing use requirements. Another significant consideration was compatibility with chosen RPHY nodes that were being tested and decided on by other teams. Taking all of these considerations into account, we built the architecture described in this paper.

The evolution of DAA infrastructure has allowed Shaw to utilize the spectrum more efficiently. It has already reduced combining losses at the plant level for installed RPDs by 6dB downstream and 3dB upstream. We have launched our first RPD-only hub site with no video RF combining. Once networking was present on-site, Shaw was able to turn this hub site up and achieve full RPD functionality in a matter of hours, due to the reduced RF combining.

### 3. Principal and Auxiliary Core Overview

When creating the initial architecture, the legacy video architecture was approached with the following three principles:

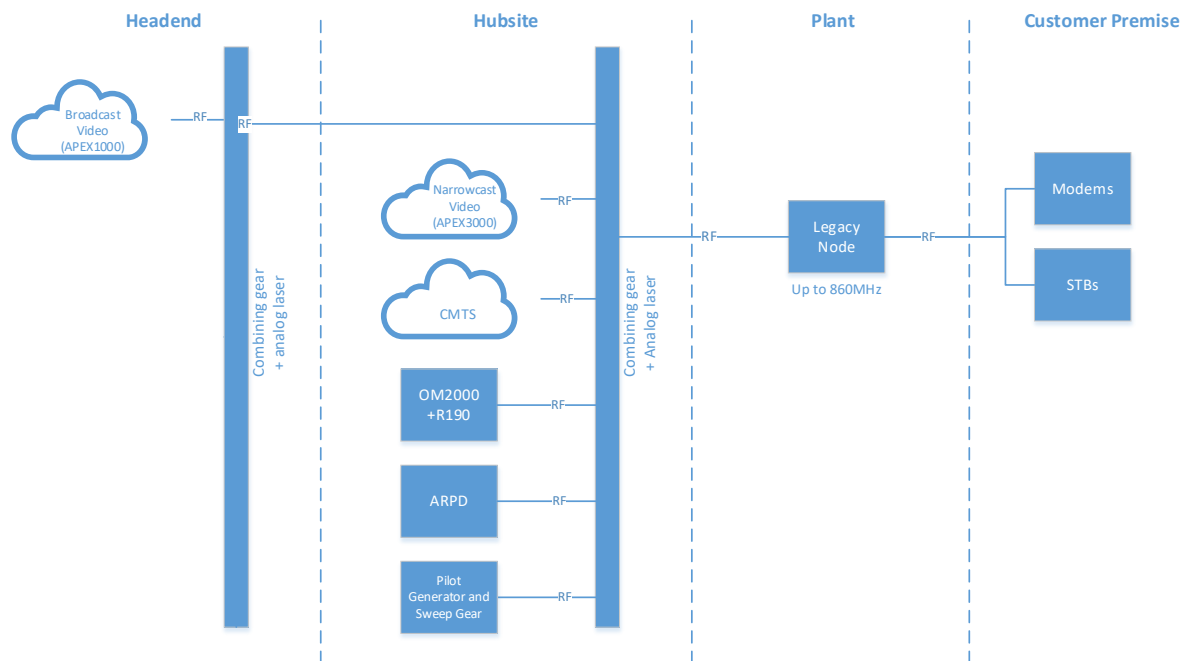
- Scalability
- Compatibility with current legacy systems
- Difficulty of automation

Having a separate auxiliary core in each region for legacy video services behind a principal core is the solution that most closely meets Shaw's current and future video needs. The auxiliary core is a secondary GCP controller core tasked explicitly with handling video RF and channel map configurations after handoff from principal core. The principal core is a dedicated GCP controller that handles initial configuration and handoff to auxiliary cores as well as all DOCSIS frequency configuration and control. This design's substantial benefit is the Cable Modem Termination System (CMTS), and video can be scaled separately depending on customer and business requirements. In the future, because of the modular architecture, the video auxiliary core can be eliminated once the transition to Internet Protocol television (IPTV) is complete with no impact to the principal core, aside from the removal of a few global settings on the CMTS. The video auxiliary core is lightweight, scalable, and relatively cost-effective compared to utilizing a CMTS video core in each region and affords a more finely tuned automation schema.

### 4. Current State of Legacy Video (Pre-DAA)

Using a headend/hub site architecture has various components distributed across the network. Most of the video components and RF combining happens in hub site facilities, which may have space and power limitations. Linear broadcast services are processed in core sites, i.e., a headend, and then combined in hub sites with narrowcast Video on Demand (VOD) and DOCSIS services to produce the full RF line up. Scaling activities and node splits can increase complexity when combining occurs in the hub site, potentially forcing the need for additional hardware.

An example of Shaw’s original legacy hub site architecture is displayed in the following diagram:



**Figure 1 Legacy Hub Site Video Architecture**

With this engineering, linear video is acquired at various signal acquisition facilities (SAF) and made available to video multiplexers—Digital Content Managers (DCMs) and edge Quadrature Amplitude Modulation (QAM) CommScope APEX1000s. Here it is distributed on Shaw’s video wide area network (WAN) over an internet protocol (IP) multicast format.

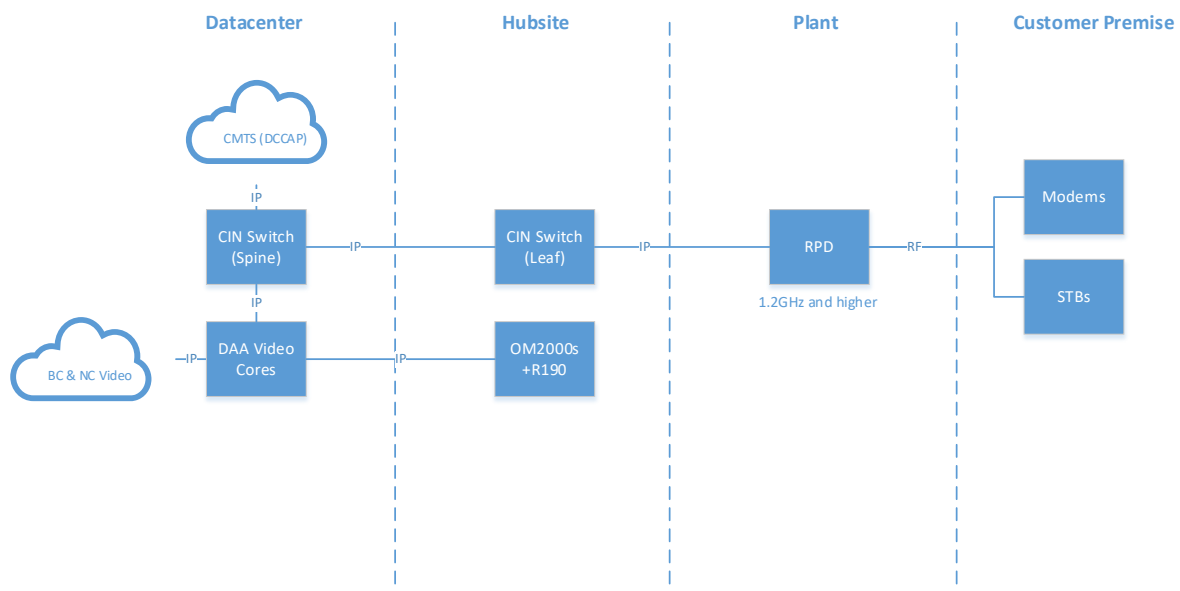
Once multiplexed into Multi-Program Transport Streams (MPTS), edge QAMs in the headend subscribe to these video multicasts and modulate the signals to video QAM carriers over RF. This broadcast RF spectrum leaving the headend is then transported to attached hub sites, where the broadcast video gets combined with local narrowcast services (VOD), and put on the plant.

## 5. DAA Technology and Components

The DAA video core infrastructure is responsible for the aggregation and encryption of video services, and the configuration of all video, OOB, and SCTE-55 communication for DAA RPDs. In collaboration with CommScope, it was possible to lay a solid foundation for a genuinely scalable DAA video architecture that creates a reliable infrastructure for rapid expansion without incurring technical debt.

Considering the long-term goal of DAA and that DAA requires the same interactive components of a legacy node (RADD [Remote Addressable Danis/DLS Downloader], OM [Out-of-Band Modulator], ARPD [Advanced Return Path Demodulator], etc.), it was decided not to create net-new. Instead, Shaw opted to utilize existing equipment already in the hub site and leverage it in the DAA solution. Each net-new and replacement DAA video node is attached to an already existing and functioning analog node or hub site in a logical manner, therefore, enabling seamless and transparent migration from analog nodes to DAA nodes in the event of a node swap or split, providing the same valued customer experience across the entire footprint.

Housing the non-legacy component of DAA in the data centre, it is possible to reduce a hub site footprint to as small as a single Converged Interconnect Network (CIN) leaf router.



**Figure 2 DAA Data Flow**

## 5.1. VUE Solution

The Video Unified Edge (VUE) provides the network and RF information between the video sources and the RPDs. It is a highly independent scalable solution designed to replace the legacy headend video gear functioning as an auxiliary core for Shaw's DAA environment. VUE is a containerized application which can be used in both bare-metal and virtualized environments.

The following servers establish the VUE:

- Two video platform server (VPS) systems operating on separate virtual machines:
  - VUE Docker repo (VPS#1), and
  - VUE platform services (VPS#2)
- VUE application servers (pipe servers), which operate on bare metal chassis


The two VUE platform servers share the functions of managing and maintaining the VUE environment. VPS#1 deploys and controls the entire VUE OS/Docker solution and performs upgrades and maintenance. VPS#2 is responsible for the Generic Configuration Protocol (GCP) communication to the RPD through the Auxiliary Core Control Plane (ACCP) sub-system, which is in a dedicated Docker container. Because this system communicates directly with RPDs, this sub-system requires connectivity to the CIN in addition to a connection to our management network.

VUE application servers (VUE pipe servers) are responsible for the creation of Downstream External PHY Interface (DEPI) and Upstream External PHY Interface (UEPI) streams. These streams carry interactive and video MPTS multicast traffic to the RPD. In Shaw's infrastructure, each pipe on the VUE pipe server is in its own Docker container and carries a single 38.8Mbit MPTS.

Internal communication between all CommScope VUE servers (applications server and VUE platform servers) use either Secure Shell (SSH) or Application Program Interface (API) over our Operations Administration Maintenance and Provisioning (OAMP) network.

The following image displays the rack elevations for the DAA infrastructure. Each network encryptor (NE) blade on an APEX 3000-NE can output approximately 6.5G of traffic, allowing Shaw to build 42 service groups per VOD COR – a VUE pipe server designated for narrowcast, at four Annex B MPTS per service group. Annex B is a cable standard that modulates a MPEG-TS input into a QAM-256 output. Each service group is assigned four VOD carriers and four RPD nodes initially. This number can be scaled up and down to dynamically meet fluctuating capacity demands utilizing our back office and automation tools.

Each DAA rack with seven VOD pipe servers can feed 1176 DAA nodes with no physical combining involved. With the second rack added, a total of 3,528 nodes can be fed.



**Figure 3 Shaw DAA Rack Layout**

## 5.2. APEX 3000-NE

The APEX 3000-NE device is a high-density network encryptor capable of supporting and ingesting 12,288 unique program identifications (PIDs); up to 4096 multicast streams. It can support Single Program Transport Streams (SPTS) for VOD and MPTS for broadcast services. The APEX 3000-NE egresses encrypted traffic on the same physical port that it ingests unencrypted traffic, allowing for a denser installation requiring fewer physical switchports.

Features:

- Eight blades configured in an N+1 model
- 12 total 10GigE interfaces where four interfaces are the back up for the first eight interfaces
- Redundant host modules
- Redundant power supplies
- Support for User Datagram Protocol (UDP) port mapped VOD and includes support for Real Time Streaming Protocol (RTSP) provisioned VOD services
- Performs broadcast and SDV encryption
- Supports PSI generation and message insertion
- Support for SCTE-52 encryption for broadcast and VOD through Common Tier Encryption (CTE)

## 5.3. Video Topology Manager

The Video Topology Manager (VTM) is a software-defined management solution that configures and manages the video configurations in DAA. Service groups with specific settings are created in VTM for each of the video services offered and then assigned to applicable RPDs. Additionally, VTM interfaces with the Digital Addressable Controller (DAC) and receives all channel information for the region. This related information defines service group information for the RPDs configured by the regional ACCP sub-system.

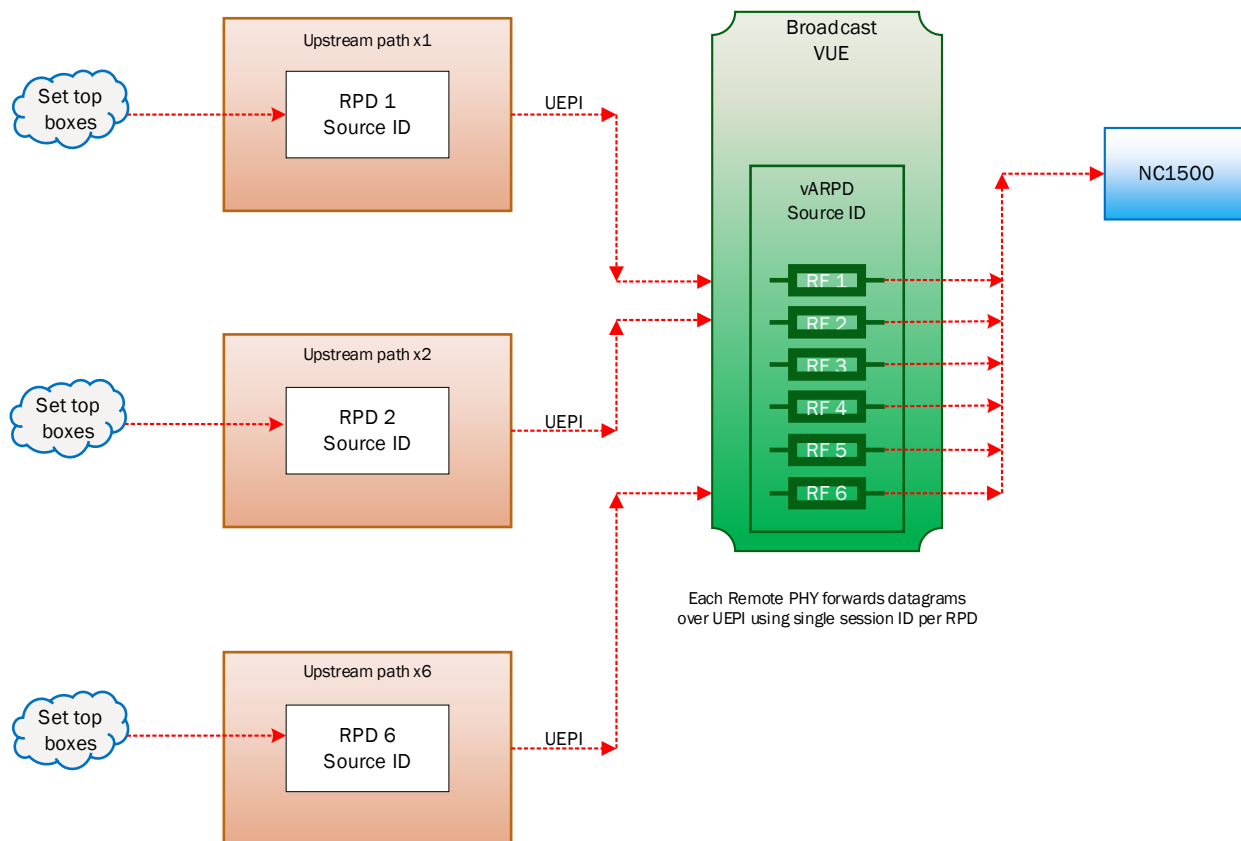
A single instance of VTM can administer the entire DAA topology. Aside from using VTM to automate, this solution can also build serving groups and modify channel maps. In the future, there will be an option to use VTM as a dashboard that enables granular troubleshooting on a per RPD basis country wide.

## 5.4. Legacy SCTE55-1 Interactive Signaling

### 5.4.1. Virtual ARPD

For interactive traffic, the Advanced Return Path Demodulator (ARPD) functionality is on the RPD. RF streams are demodulated and encapsulated using the Remote Upstream External PHY Interface (R-UEPI) standard and routed to the broadcast application server (VUE pipe server). The application server identifies the upstream traffic by the Layer Two Tunnelling Protocol (L2TP) *sessionID* and *varpdSourceid*, then routes this data to the correct port on a legacy Network Controller (NC-1500). Each Virtual Advanced Return Path Demodulator (vARPD) – one per hub site, is its own logical entity in the DAC. It is configured identically to a physical ARPD on both the legacy Network Controller and the DAC.





**Figure 4 RPD and vARPD Relationships**

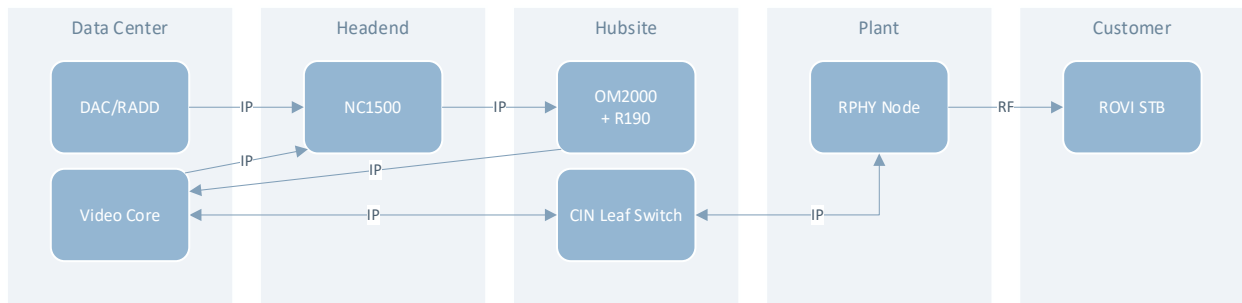
#### **5.4.2. Out-of-Band Modulator**

Downstream OOB data is still handled by an existing OM2000 located in each hub site. The out-of-band modulator (OM) sends forward data streams containing guide data, code modules, channel map information, Emergency Alert System (EAS) messaging, DAC messaging, and app network data.

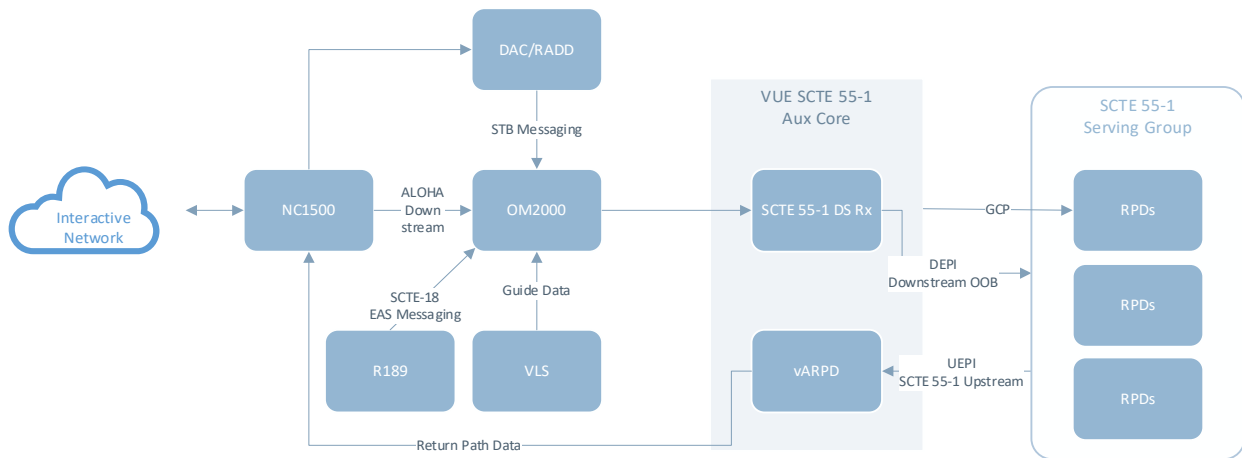
On the legacy OM2000, an IP interface is enabled pointing at the VUE application server (VUE pipe server) over the OAMP network. The OM2000 communicates with the VUE application server utilizing unicast UDP IPv4 traffic. The OM2000 packs seven, 188-byte MPEG packets into a single 1,500-byte packet and forwards to the VUE pipe server's IPv4 OAMP interface on a unique port. Each hub site OM's communication port was configured in an exclusive range per DAA region. For example, the UDP port range 4000-4100 is the Calgary region and port 4001 is specified for a particular hub site within that region and is not re-used; keeping the ports separate provides teams to discover and identify problems quickly.

### 5.4.3. NC1500

The NC1500 is a network controller that functions the same as it does in the legacy network; it works as both a DHCP server and a router which assigns IP addresses and routes upstream messages between nodes and RADDs. To complete configuration, the NC requires an IP address for each vARPD; therefore a virtual IP address is created and reserved in the Internet Protocol Access Management (IPAM) system and inserted into the NC configuration. The vARPD IP is a placeholder and not reachable on Shaw's OAMP network, so ping checks to these IP's will fail from the NC1500. Other upstream connectivity tests such as DAC refreshes are utilized instead of ping checks to verify return path functionality during troubleshooting.



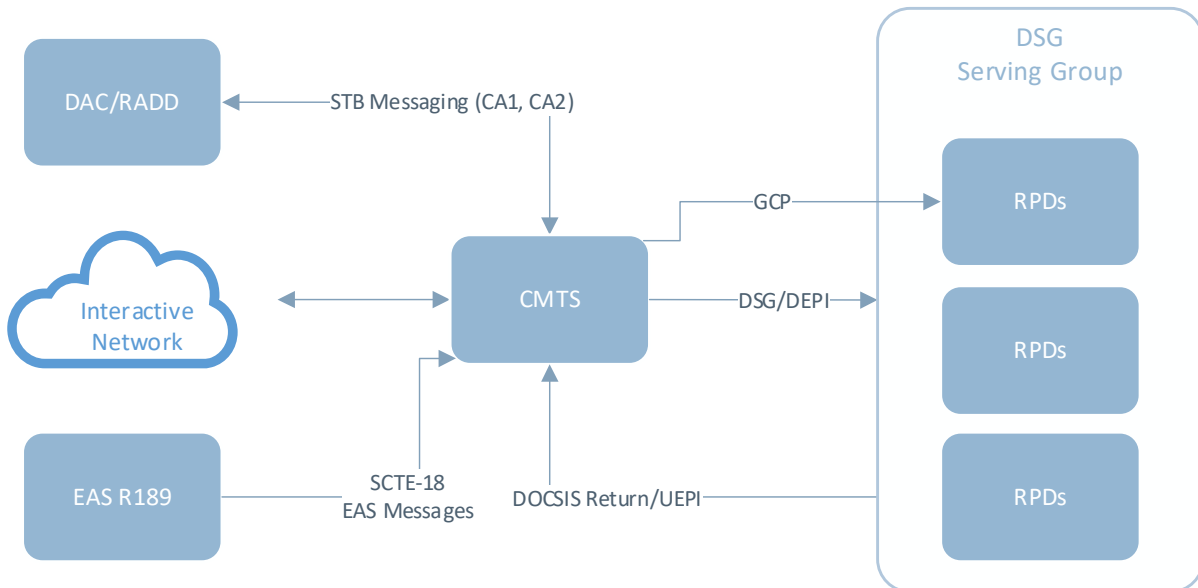
**Figure 5 DAA Interactive Network (Physical Locations)**



**Figure 6 DAA Work Flow**

## 5.5. DSG SCTE55-1 Interactive Signaling

For the DOCSIS set-top gateway (DSG) boxes, new DAA Converged Cable Access Platform (CCAP) cores have been deployed for their provisioning and control traffic. Targeting these tunnels is done at the Media Access Control (MAC) level and are configured using the existing Conditional Access (CA1 and CA2) and EAS SCTE-18 multicast streams the legacy DSG boxes already use.

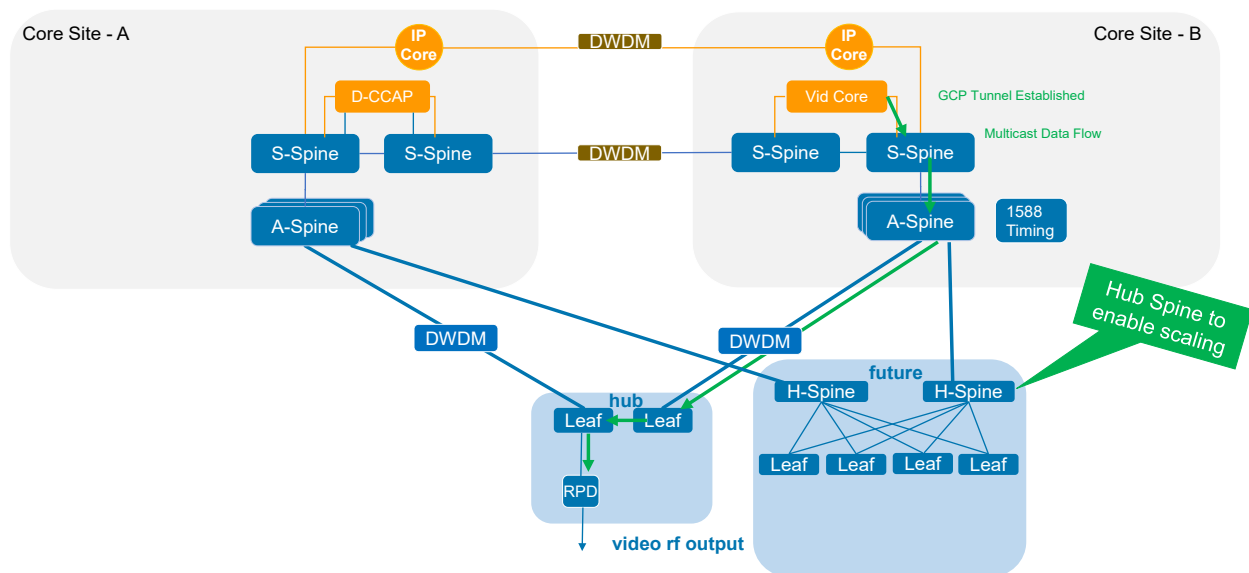


**Figure 7 DSG Interactive and Data Flow Traffic**

## 5.6. Converged Interconnect Network

CIN is a critical network that provides a redundant, highly available, scalable, and agile network solution for DAA and other Shaw services.

The CIN network leverages a spine-leaf architecture enabling a scalable DAA delivery. The spine (deployed in headend/data centres) typically has 100G interfaces and aims to maximize the fibre segments' throughput. The leaf is deployed to hub sites and aggregates traffic for multiple RPDs. All traffic carried on the CIN network is IPv6.



**Figure 8 CIN Network**

## 5.7. Timing

Shaw's DAA architecture utilizes the Precision Time Protocol (PTP), also referred to as IEEE-1588. The PTP protocol provides a method to synchronize all the devices connected to the CIN. PTP minimizes network jitter and synchronizes packet delivery across the CIN network. This is important because in current architectures, the DEPI pipes and RPDs have minimal jitter buffer. The principal core, the pipe servers, the RPDs, and any additional auxiliary cores are configured and synced to a PTP server. As with all other traffic on the CIN network, PTP is utilizing IPv6 addressing.

In Shaw's DAA design, the PTP master clock chassis syncs to the Global Positioning System (GPS) using a roof-mounted antenna, which synchronizes the two clocks deployed in each DAA region. These clocks are configured hot/hot to achieve necessary resiliency.

Currently, PTP clocks have a maximum allowance of 500 clients per clock or 1,000 clients per master clock chassis. In the future, Shaw will be transitioning to a boundary clock architecture where each RPD is assigned a more local boundary clock that syncs back to the PTP master chassis and presents itself as a PTP peer. This boundary clock architecture will allow Shaw to grow exponentially without utilizing dozens of PTP master clocks.

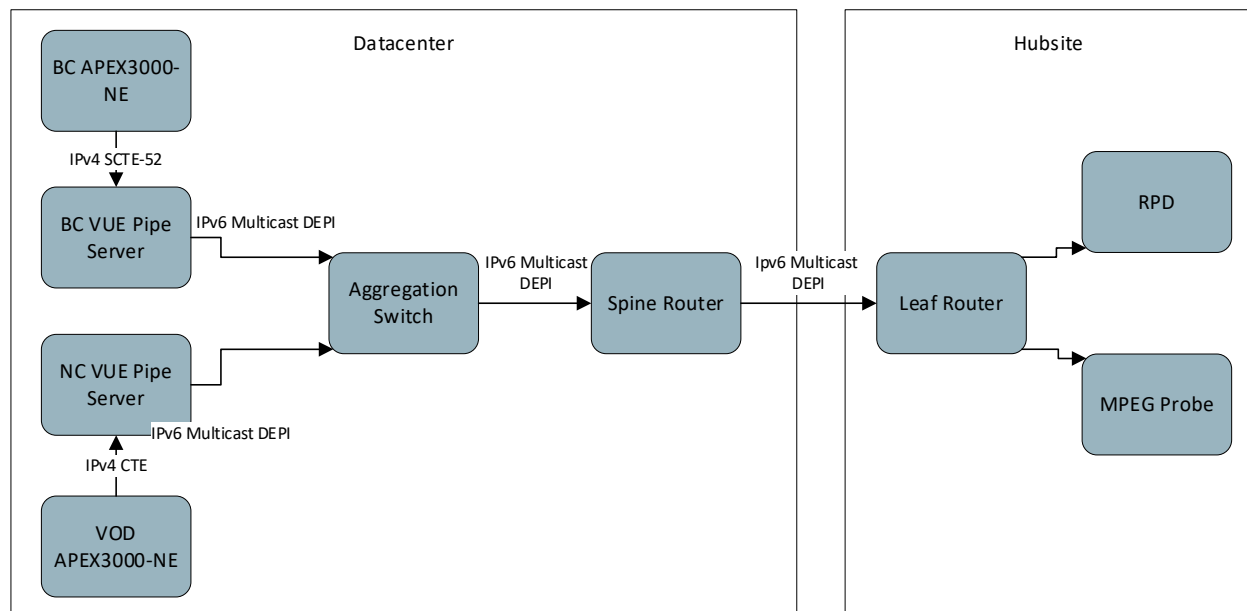
## 5.8. Remote PHY Device

The Remote PHY Devices (RPDs) replace traditional HFC nodes and assume the functionality of the edge QAMs in today's architecture. They have a network interface fed by fibre northbound, and RF interfaces southbound towards customer premises. The initiation sequence is described in the automation section [14. RPD Provisioning and Activation Process Through Automation](#).

## 6. DAA Monitoring Solution

To achieve an accurate monitoring solution, it was deemed necessary to install probes as close to the edge as possible. This accomplished two goals – end-to-end network verification and quality assurance for customer video feeds. There are probes in each region, with multiple probes in our largest region Vancouver. The probes installed in hub sites are on the same leaf routers that the RPDs are and acquire the same DEPI encapsulated multicast. The probes decapsulate multicast DEPI and then analyze the MPEG streams utilizing industry-standard methods. Shaw can currently monitor the entire broadcast channel lineup at multiple sites in real-time and trap any errors to a central dashboard. If bandwidth and licensing are concerns, narrowcast and OOB streams can be verified on an on-demand basis.

Because of DAAs regional and multicast architecture, the number of probes required can be minimized yet still maintain a high level of confidence in video quality to customers. An incumbent vendor has already integrated into the operational team's dashboard solution, so new probe outputs were added directly into an existing view. The Operations team required no further training on the monitoring solution because the probes present video quality metrics identical to the legacy MPEG probes.



**Figure 9 DAA Monitoring Architecture**

## 7. Video and VOD Encryption

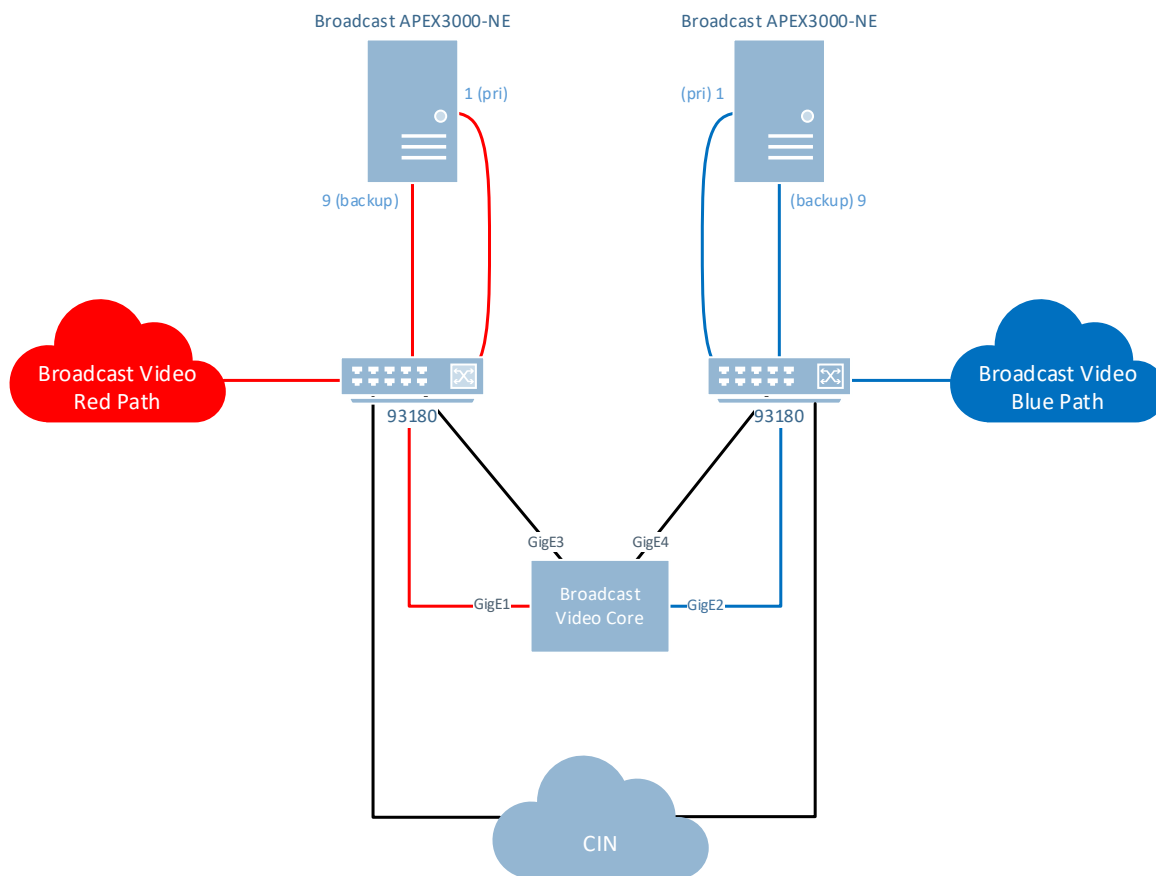
Shaw's conditional access system, the CommScope DAC, was initially configured to encrypt video services using DigiCipher II (DCII) – a CommScope proprietary encryption standard. As part of the DAA/auxiliary core project, Shaw migrated from this proprietary standard to the more open and industry-accepted alternative SCTE-52 (DVS-042). The conversion to SCTE-52 has allowed Shaw to deploy headend devices made by vendors who are not capable of licensing or using the CommScope DCII scheme.

As a sub-category of this process, all the VOD devices in production were converted to the new SCTE-52 encryption schema.

## 8. Video Broadcast Acquisition Architecture

Within Shaw, broadcast video delivery, encoding, multiplexing, and transport over the video backbone network have not changed for the DAA video solution. Instead of edge QAMs doing the encryption in the headend/data centres, video streams are encrypted by an APEX3000-NE. The NE is a product offered by CommScope based on the APEX3000 hardware with a bulk network encryptor license. Once the NE license is applied, it disables the RF ports on the chassis, receives the transport streams over IP, encrypts, and outputs over IP.

Two APEX 3000-NEs were installed at each DAA region to provide full chassis redundancy. Currently, the APEX 3000-NE cannot switch between streams. Therefore, it is impossible to utilize a single APEX 3000-NE for both primary (red) and backup (blue) broadcast video paths. Each APEX 3000-NE is connected to the primary and secondary video network. Both APEX 3000-NEs have two interfaces connected to an aggregation switch to provide link redundancy. Encrypted data streams from the APEX 3000-NEs are aggregated by the VUE pipe server, which evaluates the quality of incoming streams utilizing TR 101 290 and packet counters—in turn, selecting the best copy. The pipe server encapsulates the encrypted MPEG streams into DEPI (L2TP) tunnels.



**Figure 10 Broadcast APEX 3000-NE Data Centre Network Configuration**

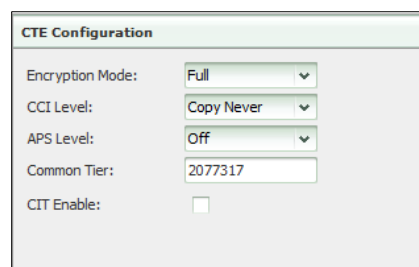
## 9. Narrowcast Architecture

Video-on-Demand (VOD) is delivered using previously existing VOD pumps that are now consolidated and centralized in the data centres in the three main regions – Calgary, Edmonton, and Vancouver. Shaw's narrowcast architecture is based on UDP mappings, whereby each RF QAM carrier is assigned a UDP port that has 21 programs (PIDs) attached to it. Program numbers 1-10 are for standard content, while programs 12 -21 are for adult content.

Currently, each QAM carrier has a unique Transport Stream Identifier (TSID) tied to it, and program 11 on each of the UDP ports is reserved for a Barker Channel embedded with TSID. This PID has been reserved in the VUE core and is automatically assigned to any port on the output UDP mappings.

A region unique TSID exists for every VOD carrier, on every VOD service group. When a VOD session begins, VOD back office (referred to as Edge Resource Manager (ERM)) maps the customer to a VOD encryptor, and calculates the next available UDP port. ERM then initiates a stream from the video server (VOD pump) to the VOD encryptor and VUE. The VOD back-office forces the STB to tune to the carrier and MPEG number mapped to the UDP port. If there is a TSID mismatch, ERM will attempt to tear down this first session and regenerate it to the correct location. This is referred to as TSID auto-correction. TSID auto-correction occurs if there is a billing system or VOD back-office configuration issue, or a physical wiring issue in the hub site. ERM is responsible for monitoring the bandwidth of the devices to ensure a VOD session can be established.

VOD encryption is also handled by APEX 3000-NE devices configured for full encryption. Each STB is authorized and entitled for VOD through our billing systems, the APEX3000-NE contains this package. Encryption is SCTE-52 and key distribution happens via the controller (DAC). The redundancy features for the APEX 3000 NE do not change from its legacy predecessor APEX 3000. No chassis redundancy has been implemented for VOD in DAA, as this mirrors currently existing VOD infrastructure on the legacy network.

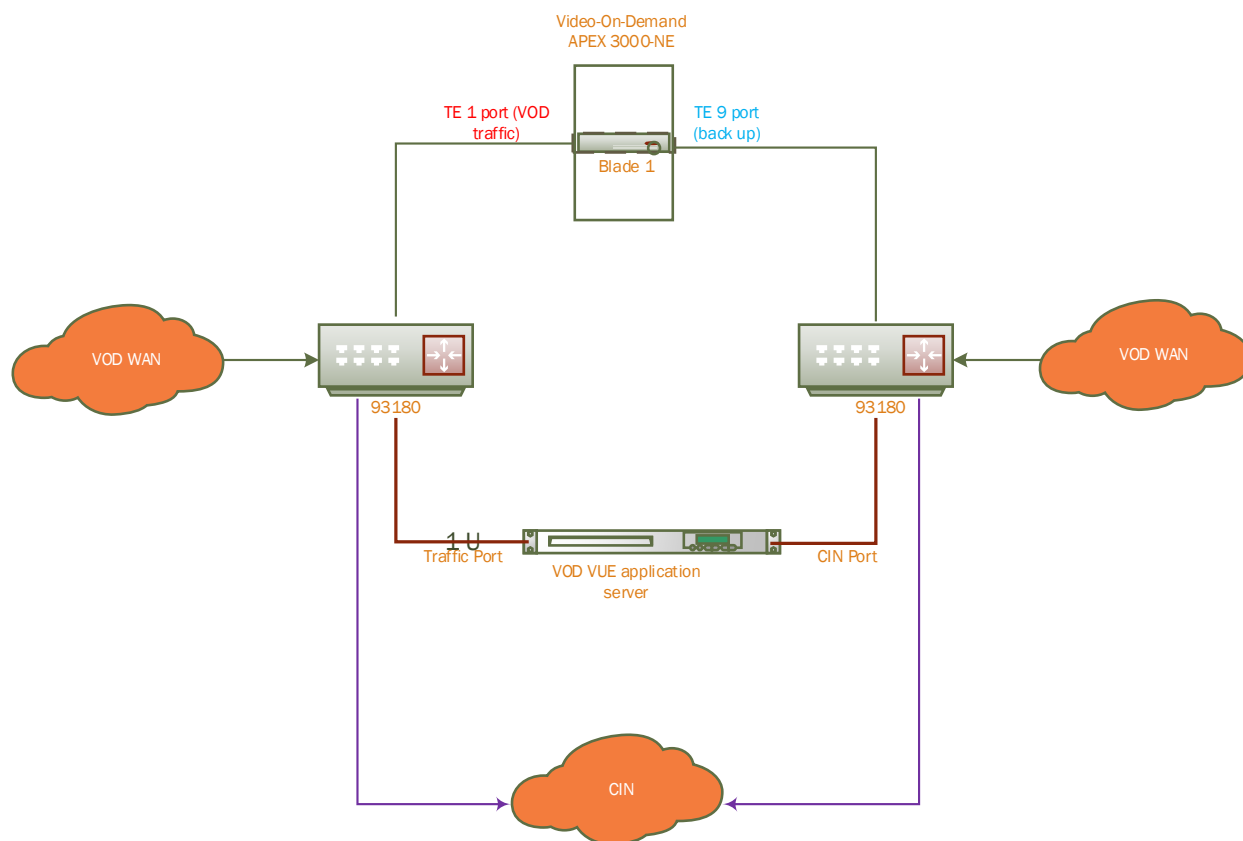


**Figure 11 NE GUI Showing VOD Encryption**

In the DAA architecture, each blade on the APEX 3000-NE is a one-to-one mapping with a VOD COR. Service groups in DAA are software-defined, and VOD sessions are streamed out over multicast. All RPD's under one single service group will subscribe to the same multicast.

A total of four interfaces are required, two each on the NE blade and VOD COR, respectively. To standardize and simplify our operations, all the four IPs selected are in the same IPv4 subnet. The output of the VOD COR propagates through the CIN network:

- Traffic interface on NE
- Virtual interface on NE
- Traffic interface on VOD COR
- Virtual interface on VOD COR



**Figure 12 Narrowcast Architecture of the APEX 3000-NE Blade and VUE Configuration**



## 10. IPv6 Addressing Standards

During the architecture process, it was decided to utilize strictly IPv6 for all addressing. This provided us with the room to logically and sustainably address RPD, cores, and pipe servers in a manner that would enable us to keep them logically separate from a regional and city perspective while also allowing for ample room for addressing. We feel that because of the number of addresses at our disposal, having separate hexets for cities and regions ease troubleshooting in the case of problems and allow the ability to be very specific when configuring network routing, assisting with keeping our network secure.

When choosing multicast addresses, a similar schema was used. Easily identifiable addresses outline OOB, local broadcast, system broadcast, and VOD just by looking at the address, while still giving us ample room to expand to as many different types of multicast as required—with no addressing space constraints.

## 11. Multicast, DEPI, and SessionID Relationships with Physical Frequencies on the RPD

RPDs assign physical frequencies directly based on the L2TP *sessionID* that DEPI tunnels deliver. In a broadcast video scenario, MPTS sessions are muxed into a DEPI (L2TP) tunnel and given the *sessionID* that points to the physical frequency that the MPTS is reconstituted onto. In our solution, *sessionID* is calculated as shown below:

```
session ID = baseMulticastSessionId + freqInMHz / 2
```

In our infrastructure, the *baseMulticastSessionId* is 0x80002000 in hex, or the integer 2147491840.

The following example calculation is utilized to determine the *sessionID* for 195 MHz in the main ACCP configuration file:

```
"channelType": "DsScQam",
"frequency": 195000000,
"multicastIp": "ff35:c531::daa:8123:1",
"tsId": 0,
"modulationType": "Qam256",
"interleaverDepth": "I128-J4",
"powerAdjust": 0
```

**freqInMHz** – 195

**sessionID** – 195 / 2 = 97

**chanObj.sessionID** – 2147491840 + 97 = 2147491937 or 0x80002061 in hex

The RPD receives the L2TP *sessionID* and reconstitutes it to 195 MHz on the plant.

Shaw has approximately 48 broadcast MPTS, and current RPD specifications only allow for 16 multicast streams. To streamline and facilitate a common architecture in all the DAA fed regions, Shaw has opted to create two unique broadcast service groups, each assigned a new multicast IPv6 address. These service groups are referred to as system and local. The system service group carries the channels common across the region, while the local service group carries channels specific to the city fed by the remote PHY nodes.

Future utilization of simultaneous substitution (simsub) content may be required (currently part of the local lineup), into a different multicast group. This schema remains within the RPDs 16 multicast session limit and allows for growth in the future.

## **12. Regional Video Realities and Solutions**

Unfortunately, the method in which Shaw rolled out plant infrastructure via acquisitions and net-new builds did not adhere to any physical frequency standard regarding the transport stream to the Electronic Industries Alliance (EIA) channel. Physical frequencies are created on the RPD; therefore, the two solutions that came forth were to either significantly increase the broadcast auxiliary core hardware and network capacity or align all the hub sites to a single physical frequency/MPTS map.

Currently, Shaw is working on aligning system transport streams (TS)/physical channels across all regional markets. Once standardization is achieved, Shaw can carry non-simsub channels using a single multicast stream, utilizing one or more secondary MPTS DEPI feeds for any simsub and local content on a per-market basis. This process will save substantial traffic across the CIN network, and a non-trivial amount of server chassis as the bulk of our video traffic is not simsub.

Due to the distance and jitter limitations of DEPI, smaller micro-cores were designed to feed smaller regional sites that are otherwise outside of the regional pipe server footprint. As long as a hub site has routes to the RED/BLUE video WAN, a single APEX-3000 NE with four raptor blades can be utilized, and three initial VUE pipe servers to feed DEPI narrowcast and broadcast to customers supplied in that remote market. This maintains GCP signalling from the regional data centres as GCP is not latency and jitter sensitive, thereby alleviating any distance concerns for DEPI video. This also reduces traffic heading across the backbone between markets and services surrounding areas with DAA without building an entirely new full-scale DAA core.

## **13. Shaw Automation Strategy**

Manually provisioning video services is time-consuming, error-prone, and a repetitive process. Shaw designed the DAA provisioning system with many small, discrete systems to handle specific tasks, rather than large systems that perform many tasks. This is a significant change from many of our past designs. Each domain (e.g. video, CCAP, CIN, DHCP) manages their provisioning components, and these systems primarily use HTTP/Representational State Transfer (REST) and SOAP as a means of communication.

Video automation has fulfilled the long-term objectives to:

1. Deploy data-and-software-driven automation to support legacy video services via RPDs.
2. Enable Shaw to deploy RPDs rapidly to meet our capacity needs to continue to offer excellent service to our customers.
3. Provide a framework that can be used to automate other CCAP components as part of any node change.

Shaw uses an in-house video automation system named VADR (Video Activation of DAA RPDs). VADR is a new system that handles the video provisioning and de-provisioning requests from Service Director (SD). Additionally, VADR supports a VOD dashboard, that assigns and re-allocates VOD service groups dynamically for RPD's depending upon their usage and the total number of nodes attached to them.

Due to the nature of DAA, multiple tools had to be either created or adapted to be functional in an RPD environment. Presently, VADR interacts with the following systems in the Shaw network to provision a node successfully. Each of these systems is either net-new or are modified to function with VADR:

- Service Director (SD)
- Resource Inventory (RI)
- VOD Central Database (CDB)
- Video Topology Manager (VTM)

### 13.1. Service Director

Service Director (SD) is a system that performs provisioning and de-provisioning of the Remote PHY Devices (RPDs). It interfaces with multiple network elements and Operations Support Systems (OSS) to complete the provisioning and de-provisioning processes.

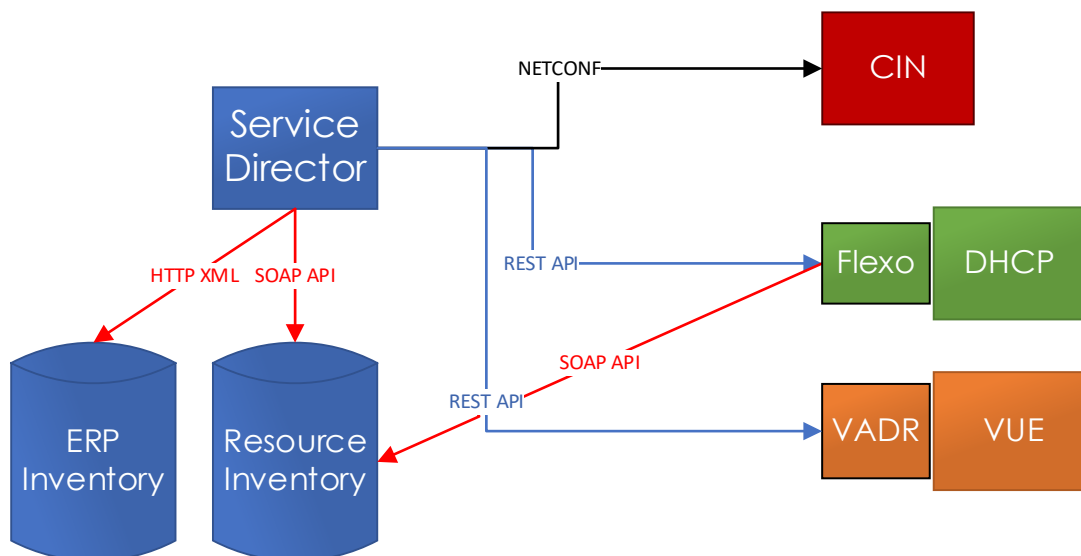
The following two items are the OSS network elements that SD interacts with to extract information and passes to the VADR for RPD provisioning.

- **Asset Inventory System (ERP):** This system provides the SD with the MAC address, serial numbers, and other RPD device-related information. SD uses the REST interface to query the ERP database for RPD parameters.
- **Resource Inventory (RI):** Details the CIN, RPDs provisioned, and other information such as IPv6 prefixes on the network.

SD interacts with multiple network provisioning elements, using the OSS components mentioned above to activate the RPD:

- **BEANS** – Automated provisioning system used to configure distributed CCAP (D-CCAP).
- **CIN (NETCONF)** – Protocol used to interact with leaf routers (hub sites), aggregation spines (headends), and super spines (data centres).
- **DHCP/FLEXO** – Automation tool used to assign IPv6 addresses to the RPDs. The SD interacts with FLEXO using the REST interface.
- **VADR/VUE** – Composes the RPDs for video configuration by receiving the provisioning and de-provisioning requests from the SD. Interaction with VADR is enabled utilizing REST calls.

The following diagram displays the SD DAA interactive map showing its interaction with multiple OSS and network elements:



**Figure 13 SD DAA Interactive Map**

## 13.2. Resource Inventory

Resource Inventory (RI) provides information about the broadcast and OOB service groups. Currently, in RI, each region is classified as a Regional Video Centre (RVC), which encompasses the region's global configurations. Each hub site under these RVC's is assigned a local and OOB service group depending upon the simsub requirements for the hub site. In Shaw's case, due to the possibility that a hub site may have a different local serving group configuration, we have enabled RI to overwrite RVC configurations with hub site configurations if a hub site configuration exists. Otherwise, the RVC configuration is passed to the RPD.

As Shaw improves the automation processes further, RI will be configured to be the source of truth for VOD COR and APEX 3000-NE configurations for VADR.

## 13.3. VOD Central Database

The VOD Central Database (CDB) is one of the critical elements of the ERM infrastructure that assigns the RPDs to a specific VOD service group and its associated multicast based on the information received from VADR. It is the central location that keeps track of and assigns return paths to each new RPD that is activated in the field.

The VOD CDB is also responsible for creating and adding VOD APEX 3000-NE blades into the VOD database tables and its correlating VOD COR.

### 13.4. Automating the Video Topology Manager

VTM is a control plane and RPD configuration system that presents a standardized API into the VUE infrastructure. Shaw uses VTM to create and modify broadcast and narrowcast serving groups, administer RPDs, and monitor and modify regional VUE solutions.

VADR communicates with VTM to configure RPDs for video. VADR collects and interprets data from SD, RI, and VOD CDB then proceeds to set configurations utilizing the VTM API. During configuration, VADR adds the RPD with the proper MAC address and relevant serving group information into VTM. VTM then takes this information and configures the VUE ACCP server itself. This allows for a highly flexible and scalable single configuration endpoint allowing for ease of troubleshooting and administration over multiple regions and a large amount of RPDs.

## 14. RPD Provisioning and Activation Process Through Automation

Due to Shaw's efforts with automation and RI, once the hub site parameters are configured into the RI and CDB systems, the bulk of RPD configuration happens seamlessly at activation.

The following steps describe an RPD node turnup:

1. A technician arrives at a location and begins wiring the node physically.
2. The technician provides the MAC address of the RPD to the implementation team.
3. The Implementation team kicks off the SD provisioning process, which bridges each of the following services, provisioning and activating the RPD on the fly:
  - BEANS configures the principal core
  - SDs NETCONF service provisions the CIN devices
  - VADR sets the auxiliary video core via VTM
4. At this point, the RPD is considered provisioned. On completion of the physical wiring, the technician can power up the RPD.

When an RPD is first connected and powered up, it receives IP addressing information and the address of the principal core through Dynamic Host Control Protocol (DHCP). Once the communication with D-CCAP is initiated, it will receive the following configuration data from the device:

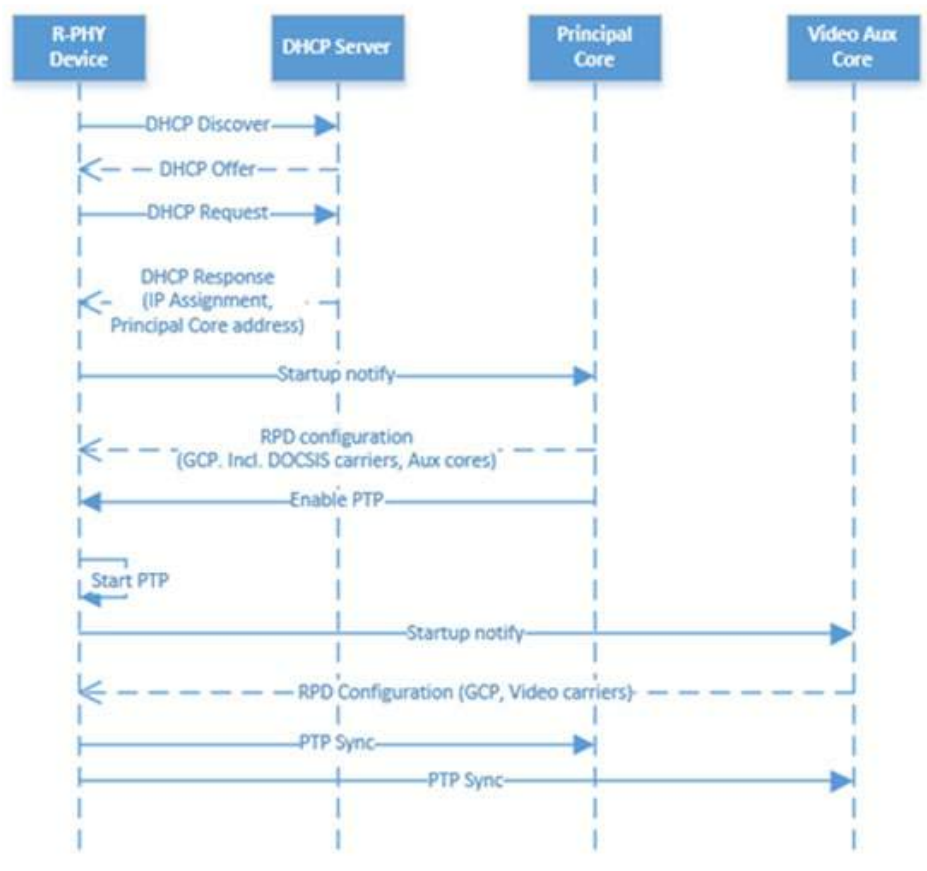
- A command to enable its PTP service.
- Network information on how to locate a PTP server.
- DEPI/UEPI and RF information (frequency, modulation, etc.), for the enablement of DOCSIS carriers.
- The address(es) of auxiliary core(s).

One of the auxiliary cores is the VUE auxiliary core. The RPD will send a startup-notification containing the RPDs MAC address to the auxiliary core, which will send the video DEPI and UEPI configuration for the RPD based upon the configuration the ACCP controller received from VTM. After the initial setup is complete, then the RPD will attempt sync with the PTP server.

The ACCP controller (VUE controller) requests a PTP sync status from the RPD after configuring all the channels and pseudowires. Once the RPD notifies VUE that PTP sync status is true, VUE then activates the previously configured channels by setting RF Mute to *Disable*, and the Admin state to *Up*. After activating all channels, VUE sets the RPD state to *Operational*. If the RPD has not yet achieved PTP sync, the channels being configured are set up in an RF Mute state, and the Admin state is set to *Down*.

If the VTM later updates channels on the RPD, the RPD has achieved PTP sync, and has been made operational, VUE will set up these channels with all of the configuration parameters; RF Mute state to *False*, and Admin state to *Up*, all at the same time.

Once the PTP sync message is successfully forwarded to the auxiliary core, all configured video RF channels are unmuted.



**Figure 14 Activation Sequence Diagram During the RPD Activation Process**

## 15. Lessons Learned

Architecting and integrating DAA with legacy video is a challenging task, and as an organization, many lessons were learned in architecting this solution.

During the architecture planning process, it was decided to utilize IPv6 for all RPD addressing. This enabled efficient use of the address space for RPD's, CIN network, and auxiliary core on a per region and hub site level while also allowing ample room for scaling out our DAA infrastructure. The sheer amount of IPv6 addresses has empowered us to have separate hexets for our regional video cores, which ease and assist with troubleshooting our network during outages. With the flexibility that IPv6 multicast addressing offers us, we decided to choose an addressing scheme that quickly identifies our out of band (SCTE-55), local, system and VOD service groups just by looking at the addresses.

We underestimated the bandwidth and multicast requirements required to service a footprint that had hundreds of unique TS/EIA frequency variants between each of our markets. Even though our channel maps reflected the same virtual channels between respective markets, this did not necessarily mean that the virtual channels resided on the same physical frequency. To be as economical as possible, we decided to align all of our TS/EIA schemas region-wide, in each region. This cut down on the duplicate multicast and the number of servers required to serve the duplicate multicast to separate markets, where the only difference is what EIA frequency the service resides on.

Early in our design phase, we came across specific markets built in a way that made sense at the time; however, as business needs or the market changed, the design and buildout of the market did not keep up. In these markets, specifically, there were conditional access requirements that differed in conditions to adjacent cities to roll out DAA video to the market. For these use cases, we had to either move the entire city between DACs or add a net-new set of DAA video pipe servers just for that city to adhere to our architectural requirements for DAA. As we advance, we take a much more holistic approach to conditional access architectures and how they interplay both now and in the future. Early in the project, the idea was that DAA should use a net-new DAC. In hindsight, it would have cost a little more, but as DAA technology improves, it may have proven to be more elegant.

Due to our sizeable footprint and the long distances between regional pipe servers and some of our more distant markets, we were outside the nominal range for DEPI. What we came up with, instead of having to create and license net new regions for smaller centers, was to create a "mini region," which would include pipe servers, but rely on the larger regional data centers for GCP command and control, as well as encryption. To facilitate this, we choose a central location for these pipe servers with a good network and a properly hardened datacentre facility and utilize a single "mini region" to feed as many small markets as possible. An excellent example of this would be our interior BC region, which has a large number of low population towns nearby, and good fibre availability, yet too much distance from a DEPI perspective to backhaul video from a fully-fledged datacentre.

From an automation perspective, early on in the project, we realized that our resource inventory system was inadequate to support our DAA initiative fully. Next was months of design and decision meetings, culminating in a multi-phased rollout to achieve full compatibility with automation tool requirements. Two of the major obstacles we had were the initial implementation of IPv6 addressing into our resource inventory system and specifying serving group, out-of-band, and VOD configurations to our automation system in a granular, sustainable, and information-complete way. As we advance with RI, we now put as much information as we have available in as specific and granular a format as possible, even if we are not currently using all the data. There may come a time when we may need to automate against it.

These are some of the more valuable lessons that we have learned while designing and implementing this solution, and we continue to learn and grow from our work in this technology daily. With features being added all the time, our architecture continues to shift and mature. The final lesson we have learned from this endeavour is to design the architecture to be modular. This avoids rip-and-replace, and duplicate efforts wherever possible.

## **16. Conclusion**

As Shaw continues to embrace next-generation technologies, we feel our architecture is adaptable and scalable enough to maintain pace with any future considerations. Because of our automation efforts, we can run two legacy video infrastructures side-by-side with minimal impact, and no new resource requirements for operational teams. As we move towards an IP-only video solution, we can systematically and remotely free up plant spectrum for CMTS, until finally completing our transition to all-IP. Our auxiliary core legacy video architecture is an essential tool in enabling Shaw's network evolution going forward.



# Terms and Abbreviations

| Term    | Definition                                                                                                                                                                                                                                                    |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ACCP    | Auxiliary Core Control Plane                                                                                                                                                                                                                                  |
| Annex B | A cable standard that modulates an MPEG-TS input into a QAM-256 output                                                                                                                                                                                        |
| API     | Application Program Interface                                                                                                                                                                                                                                 |
| APS     | Analog Protection System                                                                                                                                                                                                                                      |
| ARPD    | Advanced Return Path Demodulator                                                                                                                                                                                                                              |
| CA1     | Conditional Access 1                                                                                                                                                                                                                                          |
| CA2     | Conditional Access 2                                                                                                                                                                                                                                          |
| CCAP    | Converged Cable Access Platform                                                                                                                                                                                                                               |
| CCI     | Copy Control Information                                                                                                                                                                                                                                      |
| CDB     | Central Database                                                                                                                                                                                                                                              |
| CIN     | Converged Interconnect Network                                                                                                                                                                                                                                |
| CIT     | Constrained Image Trigger                                                                                                                                                                                                                                     |
| CMTS    | Cable Modem Termination System                                                                                                                                                                                                                                |
| COTS    | Commercial Off-The-Shelf                                                                                                                                                                                                                                      |
| CPE     | Customer Premise Equipment                                                                                                                                                                                                                                    |
| CTE     | Common Tier Encryption                                                                                                                                                                                                                                        |
| D-CCAP  | Distributed Converged Cable Access Platform                                                                                                                                                                                                                   |
| DAA     | Digital Access Architecture. Evolution and decentralization of how data centres, headends, hub sites, and nodes feed video and data services to customers' homes.                                                                                             |
| DAC     | Digital Addressable Controller                                                                                                                                                                                                                                |
| DCII    | DigiCipher II                                                                                                                                                                                                                                                 |
| DCM     | Digital Content Manager                                                                                                                                                                                                                                       |
| DCT     | Digital Cable Terminal                                                                                                                                                                                                                                        |
| DEPI    | Downstream External PHY Interface                                                                                                                                                                                                                             |
| DHCP    | Dynamic Host Control Protocol                                                                                                                                                                                                                                 |
| Docker  | A tool designed to make it easier to create, deploy, and run applications by using containers. Containers allow a developer to package up an application with all the parts it needs, such as libraries and other dependencies, and deploy it as one package. |
| DOCSIS  | Data-Over-Cable Service Interface Specifications. An international telecommunications standard that allows for the addition of high-bandwidth data transfer to an existing coaxial cable TV system.                                                           |
| DSG     | DOCSIS Set-top Gateway                                                                                                                                                                                                                                        |
| DWDM    | Dense Wavelength Division Multiplexing                                                                                                                                                                                                                        |
| EAS     | Emergency Alert System                                                                                                                                                                                                                                        |
| ECM     | Embedded Cable Modem                                                                                                                                                                                                                                          |
| EIA     | Electronic Industries Alliance                                                                                                                                                                                                                                |
| ERM     | Edge Resource Manager                                                                                                                                                                                                                                         |
| ERP     | Enterprise Resource Planning                                                                                                                                                                                                                                  |
| FLEXO   | Control plane for OSS Dynamic Host Control Protocol systems                                                                                                                                                                                                   |
| GCP     | Generic Configuration Protocol                                                                                                                                                                                                                                |
| GigE    | Gigabit Ethernet (GbE or 1 GigE) is the term applied to transmitting Ethernet frames at a rate of a gigabit per second (1 billion bits per second). The most popular variant 1000BASE-T is defined by the IEEE 802.3ab standard.                              |
| GPS     | Global Positioning System                                                                                                                                                                                                                                     |
| HFC     | Hyper Fibre Coaxial                                                                                                                                                                                                                                           |

| <b>Term</b> | <b>Definition</b>                                                                                                                                                                                                                  |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HTTP        | Hypertext Transfer Protocol                                                                                                                                                                                                        |
| IP          | Internet Protocol                                                                                                                                                                                                                  |
| IPAM        | Internet Protocol Access Management                                                                                                                                                                                                |
| IPTV        | Internet Protocol Television                                                                                                                                                                                                       |
| IPv4        | Internet Protocol version 4                                                                                                                                                                                                        |
| IPv6        | Internet Protocol version 6                                                                                                                                                                                                        |
| L2TP        | Layer Two Tunneling Protocol                                                                                                                                                                                                       |
| MAC         | Media Access Control                                                                                                                                                                                                               |
| Mbps        | Megabits Per Second                                                                                                                                                                                                                |
| MPEG        | Moving Picture Experts Group                                                                                                                                                                                                       |
| MPTS        | Multiple Program Transport Stream                                                                                                                                                                                                  |
| NC          | Network Controller                                                                                                                                                                                                                 |
| NE          | Network Encryptor                                                                                                                                                                                                                  |
| NETCONF     | A protocol defined by the IETF to install, manipulate, and delete the configuration of network devices.                                                                                                                            |
| OAMP        | Operate Administer Maintain Provision                                                                                                                                                                                              |
| OM          | Out-of-Band Modulator                                                                                                                                                                                                              |
| OOB         | Out-of-Band                                                                                                                                                                                                                        |
| OS          | Operating System                                                                                                                                                                                                                   |
| OSS         | Operations Support Systems                                                                                                                                                                                                         |
| PHY         | Physical Layer                                                                                                                                                                                                                     |
| PID         | Program Identification                                                                                                                                                                                                             |
| PTP         | Precision Time Protocol. Also referred to as IEEE-1588.                                                                                                                                                                            |
| QAM         | Quadrature Amplitude Modulation                                                                                                                                                                                                    |
| R-UEPI      | Remote Upstream External PHY Interface                                                                                                                                                                                             |
| RADD        | Remote Addressable Danis/DLS Downloader                                                                                                                                                                                            |
| RDS         | Rights Data Server                                                                                                                                                                                                                 |
| REST        | Representational State Transfer (RESTful API). API built using the rules of representational state transfer software architecture.<br>Web protocol that utilize much less overhead than SOAP. Allows for communication beyond XML. |
| RF          | Radio Frequency                                                                                                                                                                                                                    |
| RI          | Resource Inventory                                                                                                                                                                                                                 |
| Rovi        | Rovi is a global leader in digital entertainment technology for some of the largest CE manufacturers, service providers and online, mobile, and application developers.                                                            |
| RPD         | Remote PHY (Physical) Device. A device that receives Internet, telephone and video over IP and converts it to RF.                                                                                                                  |
| RTSP        | Real Time Streaming Protocol                                                                                                                                                                                                       |
| RVC         | Regional Video Centre                                                                                                                                                                                                              |
| SAF         | Signal Acquisition Facilities                                                                                                                                                                                                      |
| SCTE-18     | Emergency Alert Messaging for Cable (ANSI J-STD-42-C) standard by the Society of Cable Telecommunication Engineers                                                                                                                 |
| SCTE-52     | Data Encryption Standard - Cipher Block Chaining Packet Encryption specification standard by the Society of Cable Telecommunication Engineers                                                                                      |
| SCTE-55     | ALOHA protocol-based standard used by General Instrument and Motorola equipment                                                                                                                                                    |
| SD          | Service Director                                                                                                                                                                                                                   |

| <b>Term</b> | <b>Definition</b>                                                                                                                                                                                                                                                                                                                                |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SimSub      | Simultaneous Substitution. Signal substitution occurs when a signal is temporarily replaced by another one airing the same program at the same time. Usually, an American signal (e.g. affiliates of ABC, NBC, CBS, Fox) is replaced with a Canadian signal. Sometimes, a Canadian signal from outside the area is replaced with a local signal. |
| SOAP        | Simple Object Access Protocol                                                                                                                                                                                                                                                                                                                    |
| SOD         | Shaw on Demand                                                                                                                                                                                                                                                                                                                                   |
| SPTS        | Single Program Transport Streams                                                                                                                                                                                                                                                                                                                 |
| SSH         | Secure Shell                                                                                                                                                                                                                                                                                                                                     |
| STB         | Set-Top Box                                                                                                                                                                                                                                                                                                                                      |
| TR 101-290  | Defines measurement guidelines for monitoring MPEG Transports Streams (MPEG TS)                                                                                                                                                                                                                                                                  |
| TS          | Transport Stream                                                                                                                                                                                                                                                                                                                                 |
| TSID        | Transport Stream Identifier                                                                                                                                                                                                                                                                                                                      |
| UDP         | User Datagram Protocol                                                                                                                                                                                                                                                                                                                           |
| UEPI        | Upstream External PHY Interface                                                                                                                                                                                                                                                                                                                  |
| VADR        | Video Activation of DAA RPDs. Video Technology's service that receives provisioning/deprovisioning requests from Service Director and configures VTM using VTM APIs with data received from the various data sources. VADR also preconfigures VTM with the necessary Service Group information for VOD.                                          |
| vARPD       | Virtual Advanced Return Path Demodulator                                                                                                                                                                                                                                                                                                         |
| VOD         | Video on Demand                                                                                                                                                                                                                                                                                                                                  |
| VOD COR     | Video on Demand Core (pipe server)                                                                                                                                                                                                                                                                                                               |
| VPS         | Video Platform Server                                                                                                                                                                                                                                                                                                                            |
| VTM         | Video Topology Manager                                                                                                                                                                                                                                                                                                                           |
| VUE         | Video Unified Edge. Provides the network and RF information between the video sources and the RPDs.                                                                                                                                                                                                                                              |
| WAN         | Wide Area Network                                                                                                                                                                                                                                                                                                                                |
| XML         | Extensible Markup Language                                                                                                                                                                                                                                                                                                                       |

# References

*Remote Downstream External PHY Interface Specification*; CableLabs

*Remote PHY Specification*; CableLabs

*CCAP Architectures Technical Report*; CableLabs

*TNO-F18-D-100 Distributed Access Architecture*; Corwin Martens & Wesley Weiss - Shaw Communications

*DAA-CIN Converge Interconnect Network Solutions Design*; Shaw Communications

*Service Director Documentation*; Shaw Communications

*Architecture Definition Document – LITE*; Darren Gamble

# **To High Split & Beyond The New Frontier In Leakage Detection**

A Technical Paper prepared for SCTE•ISBE by

**Kyle Hohman**  
Network Architect  
Shaw Communications  
2728 Hopewell Place NE, Calgary, Alberta  
+1 (403) 538 5252  
Kyle.hohman@sjrb.ca

# Table of Contents

| <b>Title</b>                                             | <b>Page Number</b> |
|----------------------------------------------------------|--------------------|
| 1. Introduction.....                                     | 3                  |
| 2. Setting the Stage .....                               | 3                  |
| 3. Methodology .....                                     | 4                  |
| 3.1. Transmission Downstream .....                       | 5                  |
| 3.1.1. Notching Out Spectrum.....                        | 6                  |
| 3.1.2. Up-Conversion .....                               | 7                  |
| 3.2. Transmission Upstream.....                          | 9                  |
| 3.2.1 Generation of Carrier Waves from Cable Modem ..... | 11                 |
| 3.2.2 Utilizing Cable Modem Pilots.....                  | 11                 |
| 4. Conclusion.....                                       | 12                 |
| Abbreviations.....                                       | 13                 |
| Bibliography & References .....                          | 13                 |

## List of Figures

| <b>Title</b>                                              | <b>Page Number</b> |
|-----------------------------------------------------------|--------------------|
| Figure 1 - Transition from 85 MHz to 204 MHz.....         | 4                  |
| Figure 2 - DSB-SC Signal Block Diagram .....              | 5                  |
| Figure 3 - Conceptual Spectrum Notch .....                | 6                  |
| Figure 4 - Additional Group Delay Locations .....         | 7                  |
| Figure 5 - Up-Conversion.....                             | 8                  |
| Figure 6 - Spectrum Available For Modem Transmission..... | 9                  |
| Figure 7 - Conceptual Modem Transmission.....             | 10                 |

## List of Tables

| <b>Title</b>                                                               | <b>Page Number</b> |
|----------------------------------------------------------------------------|--------------------|
| Table 1 – Calculated Available Upstream Capacity based on Notch Width..... | 7                  |
| Table 2 – Capacity for Modem in Up-Conversion Plant .....                  | 9                  |
| Table 3 – Full Upstream Available .....                                    | 10                 |
| Table 4 – SC-QAM CW Harmonics.....                                         | 11                 |
| Table 5 – Summary of Options .....                                         | 12                 |

# 1. Introduction

As the industry explores providing higher bandwidth services, leveraging the Data-Over-Cable Service Interface Specification (DOCSIS) 4.0 and look at evolving legacy DOCSIS 3.1, operators are planning to move to 204 MHz split and higher return band splits. In order to meet the existing plant leakage regulatory requirements and ensure that operations are not interfering with licensed over the air signals, operators need to evolve how they measure and monitor cable plants for signal leakage.

This paper will discuss the options which are available to operators, balanced by what is technically possible based on the products available today and the roadmaps for the future. Two methodologies will be explored in this paper. The first method examined will be the use of the traditional downstream transmission of signals through the plant and the options that exist within the preservation of this existing form of transmission. The second method discussed will use the cable modem to generate specific signals to characterize leaks and transmitting them back through the plant.

For context, the paper will present potential deployment scenarios and the associated impact to the existing leakage programs that are being run in the industry and how these options can impact the requirements around the legacy set top box Out-of-Band (OOB) control telemetry.

With an open discussion and alignment about ways to solve these problems and a commitment from industry to build scalable and affordable solutions into future technologies, the industry can move forward and continue to meet regulatory requirements seamlessly.

## 2. Setting the Stage

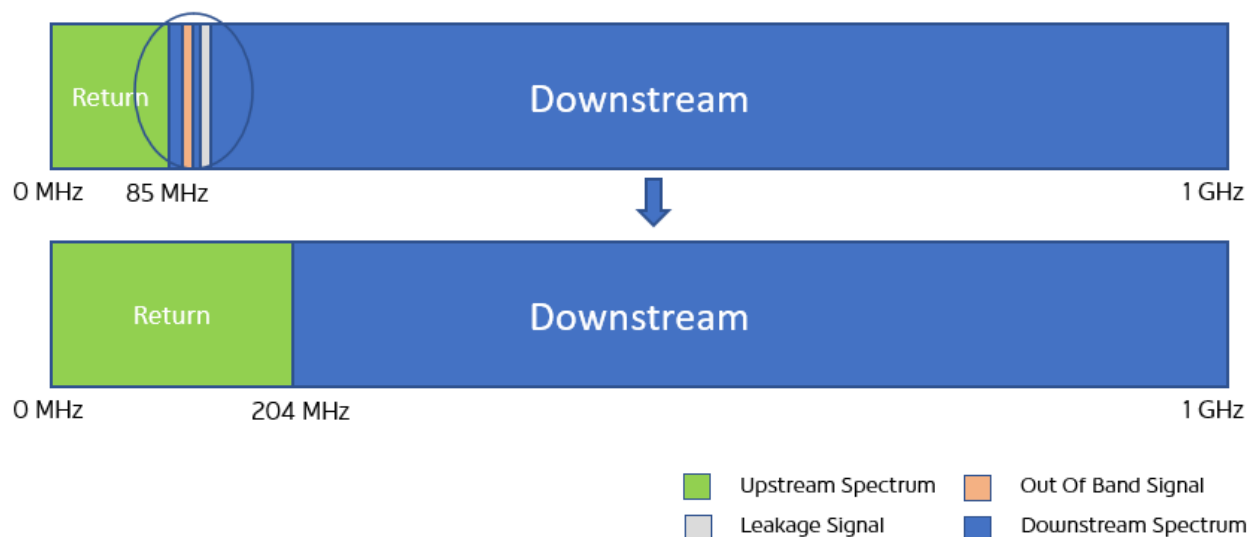
Moving to high split is a transformative network change which entails various technical and operational challenges as operators upgrade the relevant plant components. However, operators are up to this challenge because increased upstream will enable them to offer higher speed tiers and unlock additional revenue potential.

Around the world, cable companies operate in different regulatory environments. However, their foundational principals are quickly converging onto a similar framework, and the OOB requirement is agnostic to regulatory requirements. When it comes to high split, every operator is expected to reach the same end state, though different regulatory environments may impact the “how.” In addition, each company will have a unique set of business requirements, plant upgrade roadmaps and test equipment configurations that will dictate the optimal path forward.

There are two fundamental designs for the plant transmission and maintenance of leakage signals. The first is the existing paradigm is the generation of the signals in the hubsite. The second is an approach being developed around having the signals generated in the customer premises.

Operators will likely go through a transitional period where they move to provide higher upstream services, while balancing business and regulatory requirements. A conceptual, roadmap-based approach is required as operators move toward the “end state” for signal leakage detection within the aeronautical band. Throughout the discussion below there will be some simplifications made in both the graphical representation of the concepts, and the descriptions of the solutions therein. These simplifications are made to make the options accessible and clear for the reader, while not removing any of the logical or conceptual challenges that are described within the solutions.

The shift in spectrum is depicted in Figure 1 below.



**Figure 1 - Transition from 85 MHz to 204 MHz**

Due to the fundamental properties of signal transmission in the access network and customer premises, any signal trying to traverse through the access network plant in a direction opposite to the allocated spectrum (i.e. upstream signals on the downstream or vice-versa) may cause interference. These signals will also be blocked when they encounter a diplex filter in an active or passive piece of equipment. There are some creative solutions around using these features of the plant to the operator's advantage which will be explored throughout this paper.

The other consideration operators will need to address throughout this transition is their oversubscription rate, which can be described as the percentage of available capacity being offered to each customer over and above the capacity required for each customer to achieve their maximum data rate simultaneously. Upstream capacity complicates the calculation when compared to the downstream oversubscription rate because large portions of upstream traffic are in bursts and latency sensitive. The in-depth considerations of these implications will not be addressed in this paper, where the focus will remain on the transmission of different telemetry signals.


### 3. Methodology

Throughout this section, we will cover the direction of transmission of the signals through the plant. There will be the existing downstream transmission of the signals from the hubsite to the customer premises and then that existing paradigm will be inverted and the transmission of these leakage signals will be generated in the customer premises and back towards the hubsite.

There are two main types of leakage signals that will be referenced throughout the course of this paper. The first signal is a dual sideband suppressed carrier (DSB-SC) signal. This is demonstrated conceptually below and was the next step of leakage signals after switching from analog video carriers. Today this is



generated by two carrier waves (CW) that have a specific frequency offset and amplitude relative to the regular downstream transmissions on the plant.



**Figure 2 - DSB-SC Signal Block Diagram**

The second type is contained within the DOCSIS 3.1 specification, specifically the utilization of pilot carriers that exist within orthogonal frequency-division multiple access (OFDMA) known as OFDM upstream data profile (OUDP)

Tables of calculations that follow are based on the following assumptions:

- Single carrier quadrature amplitude modulated (SC-QAM) channels are 6 MHz wide with a capacity of 25 Mbps
- OFDMA channels are 1024-QAM with an expected bit rate of 8 bits/sec/Hz

These numbers and assumptions are approximations based on information available today.

### 3.1. Transmission Downstream


There are two methods that leverage existing downstream transmission signals. These methods have compromises made in terms of spectral efficiency for the newly available upstream and require a complex amplifier design that is expected to result in a higher cost for this equipment. This is balanced by two main benefits, both of which are centered around the retention of existing leakage detection test equipment and video technology platforms. The intent of this is not to provide a cost analysis but to succinctly lay out the options that exist if an operator wishes to preserve these legacy signals.

The first primary benefit is the retention of the existing leakage program that is currently being used by the operator. Depending on the size and scale of the operator and their operations, outfitting the field teams with new leakage equipment may prove to be a large undertaking from a capital standpoint. Additionally, as an industry, there is an increasing focus on the operational cost of technology changes and any time there is a change to equipment, practices, and methodology there is the introduction of an operational cost that will also need to be captured.

The second primary benefit is preserving any relevant downstream signals, specifically OOB. The removal of the OOB signal is not an option for many operators, as the number of set tops boxes that rely on only this signal and cannot use other telemetry is in the millions in many cases.

### 3.1.1. Notching Out Spectrum

When looking at keeping the existing legacy signals intact, the first thing considered was preserving a portion of the existing downstream spectrum by using notch filters for the switch to a 204 MHz return, known as high split, as pictured in Figure 3.



**Figure 3 - Conceptual Spectrum Notch**

Sacrificing any portion of the newly increased upstream spectrum, especially as operators look to reap the benefits of the newly opened return band in terms of both congestion relief and providing higher upstream tiers to customers, may seem counterproductive at first glance.

The cost of creating this notch in the upstream and sacrificing the return spectrum should be understood contextually in terms of what it means from a capacity perspective for the operator, what it means for speed tiers and potential to help capacity to help relieve congestion. The capacity calculated itself is not sufficient for determining the true costs of this spectrum sacrifice to a specific operator.


To determine the size of the notch required in the spectrum to preserve the downstream transmission of legacy signals, it is important to keep in mind that the OOB signal has some room to move within the downstream spectrum. The leakage signal may also have some room to move within the aeronautical band depending on the regulatory environment of the operator. Placing these signals as close together as possible will minimize the size of the notch required and keep the available capacity as high as possible given these constraints.

The table below details three potential notch-width scenarios based on manufacturing technology available today, and what can be reasonably expected for spectrum required to preserve the signals. A thinner notch would allow for the preservation of the OOB signal, while a wider notch could accommodate both the OOB and the leakage signals. For the purposes of this calculation, it is assumed that there would be six 6.4 MHz wide SC-QAM channels and the remaining of the spectrum would be utilized by OFDMA.

**Table 1 – Calculated Available Upstream Capacity based on Notch Width**

| Notch Width | Available Spectrum  | SC-QAM Spectrum | OFDMA Spectrum       | Projected Capacity |
|-------------|---------------------|-----------------|----------------------|--------------------|
| 20 MHz      | 5-105 & 125-204 MHz | 5-55 MHz        | 55-105 & 125-204 MHz | 1182 Mbps          |
| 30 MHz      | 5-105 & 135-204 MHz | 5-55 MHz        | 55-105 & 135-204 MHz | 1102 Mbps          |
| 40 MHz      | 5-100 & 140-204 MHz | 5-55 MHz        | 55-105 & 140-204 MHz | 1062 Mbps          |

There is another consideration for notch filters in longer cascades known as Group Delay. Group Delay is caused by a multitude of issues, which are relatively known, understood, and compensated for in deployments today. The introduction of a notch filter will create the same characteristics around the upper and lower bounds of this filter that cause group delay in upstream transmission. It is expected that an operator will end up with group delay issues on either side of notch filters in longer cascades in addition to the other group delay issue that are expected for the channel bounded by the high split filter. The spectrum that can expect to see group delay issues arise in is highlighted in Figure 4.



**Figure 4 - Additional Group Delay Locations**

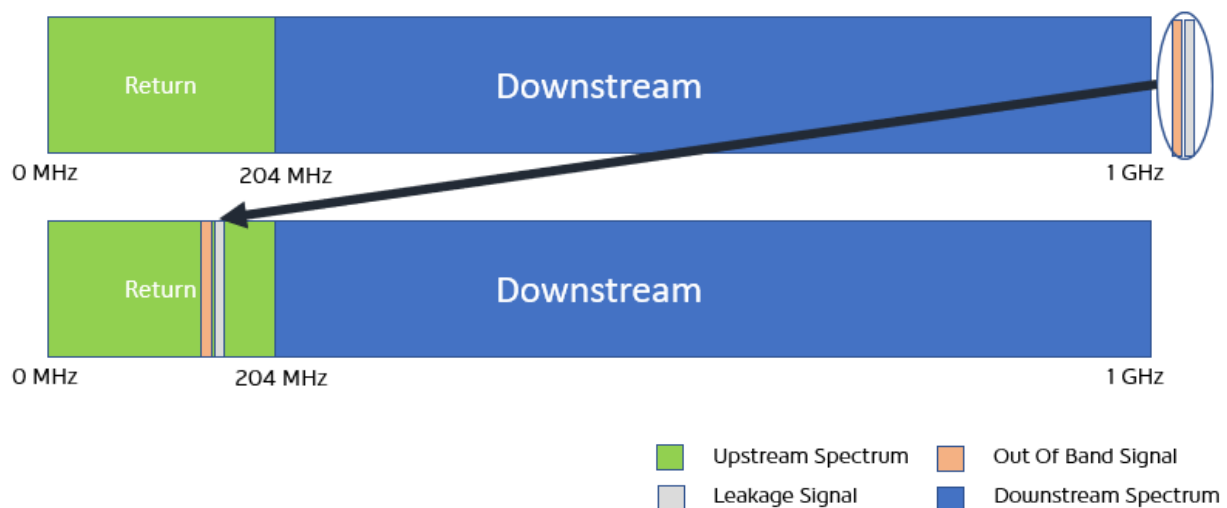
There are not currently any deployments of this solution currently available for study and evaluation, so any discussion around this expected loss in modulation for the OFDMA channels that border this filter will be based on existing upstream deployments and understanding of OFDMA. It is difficult to say whether the OFDMA channel is so robust that it can overcome the majority of the impairments introduced from group delay, or if it will degrade the whole channel to a point that leaving a small exclusion band is preferable. An in-depth analysis of the optimal size of the notch filters and exclusion channels requires future investigation and is beyond the scope of this paper.

### **3.1.2. Up-Conversion**

The second method of downstream transmission that will be discussed will be referred to as up-conversion. This concept is based on the idea of transmitting the signals which are required higher in the spectrum than where they are needed, in what would still be downstream. Once the signal reaches an active device in the plant, it will be transposed from the higher frequency and inserted as a downstream transmission in the return band on each output legs of that active. Once the signal is inserted into the return band, it begins to traverse that span of coaxial cable in a direction opposite to the other signals on that cable. It would be bound by the diplexers in the plant, where the signal would effectively be discarded since it would be trying to traverse the filter in the wrong direction. These leakage detections and OOB signals would then be regenerated on the other side of the active equipment.


For this Up-Conversion, there is additional complexity created in the amplifiers and downstream spectrum plan when compared to adding a notch filter, but up-conversion could be a more efficient use of the available return spectrum. There is also the option to handle the OOB and leakage signals differently in this deployment. The OOB will still have to be generated in the head-end or hubsite and transmitted through the plant since it is a signal that requires specific intelligence and transmission characteristics. The leakage tones however may be generated at each amplifier, or at the node, as required. This is due to the fact that although the signal is very specific, it does not require a two-way communication path or any acknowledgements from the equipment receiving the signal. The option to transmit the OOB in the downstream and convert the leakage signals as required in actives, may prove to be advantageous in the manufacturing of the actives, tuning or deployment. However, the benefits are not measurable at this time.

This is conceptually depicted in Figure 5.



**Figure 5 - Up-Conversion**

The inclusion of these signals on the return band does come with a reduced return spectrum available for transmission at the customer premises. To the modem in the home, these incoming signals will be blocked by the diplex filter within the modem and not cause interference. The spectrum that these signals are arriving into the modem on however will be occupied, and therefore unavailable for the modem to use in upstream transmission. The specific placement of these signals will likely vary in every cable operator deployment due to various interpretations of the optimal spectrum plan, manufacturer guidance and other debate, but the fact remains that there will be a small area of upstream spectrum that the modem should not be allowed to transmit on due to the fact that there will be existing signals there. This is illustrated in Figure 5.



**Figure 6 - Spectrum Available For Modem Transmission**

The size of the notch in Figure 6 is not to scale and is provided for illustrative purposes only. The amount of unavailable spectrum will vary depending on how the signals are placed and layered into the spectrum. The calculation provided in Table 2 could be considered a conservative estimate of what would be available for upstream capacity.

The group delay issue reappears in the up-conversion deployment due to the fact that it will still require a filter on the output leg of the amplifier to insert the leakage and OOB signals without further interference or ingress. The placement and function of the filter would be only on the output of the amplifier, and not the input which would lead to approximately half the group delay impact of the notch filter.

**Table 2 – Capacity for Modem in Up-Conversion Plant**

| Available Upstream Spectrum | SC-QAM Spectrum | OFDMA Spectrum       | Projected Speed |
|-----------------------------|-----------------|----------------------|-----------------|
| 5-105 & 115-204 MHz         | 5-55 MHz        | 55-105 & 115-204 MHz | 1262 Mbps       |

This allows for a more efficient usage of the upstream spectrum than the notch method and largely avoids the issues introduced by group delay. This comes at a cost of increased complexity in both sides of the spectrum deployment due to the specific configuration required for the modem, the likely requirement of a very specific upstream channel plan and finding space for the signals to be transmitted in the downstream.


### 3.2. Transmission Upstream

Leveraging the modem to generate the required leakage signals would allow for the signal to be created in every area of the plant, where there is live signal. These methods provide what is required for the leakage signal generation and having it created by the modems themselves ensures that they are in every portion of the plant that is transmitting in that spectrum. Conceptually, this is essentially turning the transmission of these signals around, where before they were going from the hub site to the customer, now they will go from the customer into the network.

While discussing the solution based around the modem transmit leakage signals, the plant being modeled is one that has removed the OOB requirement for the set top boxes. The reason for the removal of the OOB signal is not a technical requirement of the solution and it is possible to maintain the downstream transmission of this signal in the following scenarios. The reason for OOB from the discussion for the

upstream transmission scenarios are twofold. The first is that the methods of preserving the signal have been covered above, since it will need to remain a downstream signal that has a two-way communication to the hub site. The second is that removal of this signal before inverting the leakage model is programmatically aligned with the tradeoffs considered in the transmission methods, as well as conceptually aligned with transforming the plant in a manner that is easy to understand.

The measurement of these signals is in the same portion of spectrum that is currently measured and meets existing regulatory requirements in the aeronautical band without adding great complexity to the rest of the plant. Figure 6 is a high-level diagram of how this would look, with the new transmission area of the modem depicted in purple.



**Figure 7 - Conceptual Modem Transmission**

Initial findings from silicon manufacturers centered around having the modem transmit the required signals have led to the assumption that there will be minute-to-negligible sacrifices in the upstream, allowing the full spectrum to be provided as capacity as pictures in Table 2. The distribution of spectrum between SC-QAM and OFDMA channels has remained due to the fact there will be previous generations of DOCSIS in the plant for a long time to come.

**Table 3 – Full Upstream Available**

| Available Upstream Spectrum | SC-QAM Spectrum | OFDMA Spectrum | Projected Capacity |
|-----------------------------|-----------------|----------------|--------------------|
| 5-204 MHz                   | 5-55 MHz        | 55-204 MHz     | 1342 Mbps          |

The deployment of this solution has a caveat that comes from the following scenario. The Operator has taken node ABCD and upgraded it to 204 MHz. At this moment, all the existing modems, or at least the majority, only support 42/85 MHz split, and the operator may be out of compliance with a regulatory body. This is a challenge every operator will have to resolve. One suggested path involves ensuring that there is a modem installed, either in the customer premises or an environmentally hardened cabinet, that is transmitting the leakage signal back through the plant from the last tap in every run. This modem cannot distribute the required signal to all the other drops on the tap run due to the nature of coaxial distribution plant, specifically around port-to-port isolation between taps. The lack of leakage signal on every drop without a 204 MHz return modem is not immediately a problem. This is because if there are no 204 MHz return modems in an area, there is no transmission within that band that will need to be measured. The 42/85 MHz modems will continue to function as normal, and not transmit within those restricted bands. As the homes on the node become populated with 204 MHz modems that can transmit in the areas that require monitoring, those modems will be transmitting the required aeronautical band leakage signals for detection.

### 3.2.1 Generation of Carrier Waves from Cable Modem

The first option examined for the inversion of signal generation will be using the Cable Modem to generate the CW tones at spacing (refer to Figure 1 for an illustrative example) and power for the existing leakage equipment to detect. The immediate advantage in this method are clear, the cable modem can effectively generate the same signal that the leakage detection equipment has been detecting before, and at very low cost in terms of upstream spectral usage.

Using the chipset to generate the two CW tones does count towards the maximum number of channels available on the tuner. The chipset that was examined for the purpose of this analysis was the Intel Puma 7. This chipset has eight SC-QAM channels available on the tuner, which would be reduced to six available channels while having these two additional tones allocated for leakage detection. This means that there would be a limitation to the number of SC-QAM channels that could be run, and the remainder would need to be OFDMA. It is not considered to be a material cost because most operators do not run more than six SC-QAM channels today. Once the spectrum is expanded, it is much more spectrally efficient to add OFDMA channels instead of further SC-QAM channels.

During testing of this methodology where the modem chipset generates these narrow and powerful tones, some harmonic signals have been observed. The behavior of the CW tone generation and resulting side effects are detailed in Table 4.

**Table 4 – SC-QAM CW Harmonics**

| <b>CW Center Frequency</b> | <b>CW Signal Strength</b> | <b>Harmonic Location</b> | <b>Harmonic Strength</b> |
|----------------------------|---------------------------|--------------------------|--------------------------|
| 85 MHz                     | 34 dBmV                   | 215 MHz                  | -38 dBmV                 |
| 108 MHz                    | 33 dBmV                   | 192 MHz                  | 4 dBmV                   |
| 137 MHz                    | 31 dBmV                   | 163 MHz                  | 23 dBmV                  |

The creation of these harmonics may require additional regulatory, specification or investigation considerations that are beyond the scope of this paper.

### 3.2.2 Utilizing Cable Modem Pilots

The DOCSIS 3.1 standard includes protocols within OFDMA that can be used as pilots for leakage detection in the inverted leakage detection paradigm. There is a feature known as the OFDM upstream data profile (OUDP) that utilizes very specific pilot signals and although they are designed for a different purpose, there is the potential to leverage these pilot signals for leakage detection. Through optimized cable modem termination system (CMTS) scheduling, an operator may be able to remove any encumbrances to the spectrum entirely and allow a fully open return spectrum free from any notches, areas unavailable to transmit, or destructive interference.

The utilization of an existing carrier or protocol which does not add additional costs or limitations to the access network would be a huge benefit to any leakage program. Leakage has evolved from requiring whole channels, to potentially fitting in between SC-QAMs. The end goal is to use existing signals, transmitted without a penalty in bandwidth, to meet regulatory and plant hardening requirements.

There are three main areas to consider when it comes to using OUDP for leakage detection:

- The DOCSIS side requires creation of a CMTS scheduling scheme to minimize the cost to the access network. The evolution of the cable modem transmission profiles to the point where the OUDP pilot signals represent a consistent and distinct signal to be measured and profiled are still in the exploratory stages. The framework exists within the DOCSIS standards to support this initiative.
- From the test equipment perspective, measurement of these new pilot signals will be a technical challenge for equipment manufacturers. This is a new type of signal being generated, and it will have different propagation and measurement characteristics. At minimum, this will require firmware upgrades. For existing leakage test equipment, a hardware upgrade may be necessary.
- After the technical challenges are resolved on how the signal will be generated, scheduled and measured in the field, there will need to be extensive testing around the signal equivalency to existing leakage systems to ensure regulatory compliance.

## 4. Conclusion

This paper examined four solutions for leakage and OOB issues which occur in high-split and the considerations for each of these options. There are advantages and disadvantages to each, and the opportunity for other permutations of these solutions is still possible. Table 5 summarizes the proposed solutions and the main factors that have been discussed.

**Table 5 – Summary of Options**

| <b>Proposed Solution</b> | <b>Available Capacity</b> | <b>OOB Removal Required</b> | <b>Replace Test Equipment</b> |
|--------------------------|---------------------------|-----------------------------|-------------------------------|
| Notch                    | 1062 Mbps (min)           | No                          | No                            |
| Up-Conversion            | 1262 Mbps                 | No                          | No                            |
| Modem CW                 | 1342 Mbps                 | Yes*                        | No                            |
| Modem OFDMA Pilot        | 1342 Mbps                 | Yes*                        | Yes                           |

\*If OOB is left on the plant, it is still possible to move to these solutions. The capacity available will closely resemble the in the “Up-Conversion” solution.

The industry needs to align on the optimal milestones in the roadmap for evolving leakage detection. This will drive the development of solutions, as well as testing and detection of these new signals and configurations.

Once operators align on one solution, the industry will be better equipped to advocate for their needs with both regulators and suppliers. This will allow for manufacturing efficiencies and lower maintenance and upgrading costs for plants.

The size and scale constraints that come with a physical cable plant drive the need for a thoughtful, balanced approach to change. This can be a delicate act with regulatory requirements, operational requirements, business demands and technology changes. Evolution of leakage detection methods must happen concurrently with the evolution of the cable plant; it cannot be treated like an afterthought.

Creating an open and transparent dialogue around the move to high-split creates an excellent opportunity for operators to transition smoothly to high-split and continue to increase their offerings to customers, especially at a time when upstream is becoming more important than ever.



## Abbreviations

|        |                                                 |
|--------|-------------------------------------------------|
| OOB    | out of band                                     |
| DOCSIS | Data Over Cable Service Interface Specification |
| MHz    | forward error correction                        |
| CW     | carrier wave                                    |
| OU DP  | OFDM Upstream Data Profile                      |
| Hz     | hertz                                           |
| MHz    | megahertz                                       |
| GHz    | gigahertz                                       |
| OFDMA  | International Society of Broadband Experts      |
| DSB-SC | Dual Sideband Suppressed Carrier                |
| CMTS   | Cable Modem Termination System                  |
| SCTE   | Society of Cable Telecommunications Engineers   |

## Bibliography & References

*Data-Over-Cable Service Interface Specifications DOCSIS® 3.1*, Cablelabs

*Leakage detection using SCQAM channels*, Intel

# **Improving The Latency Of An MSO Network For Gaming And Real Time Applications**

A Technical Paper prepared for SCTE•ISBE by

**Colin Dearborn**

Sr. Network Architect II  
Shaw Cablesystems G.P.  
2728 Hopewell Place. NE, Calgary, AB, T1Y 7J7  
(403) 538-5297  
colin.dearborn@sjrb.ca

# Table of Contents

| Title                                                                        | Page Number |
|------------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                         | 4           |
| 2. Current Methods For Improving Latency and/or Jitter .....                 | 7           |
| 2.1. In-Home Technologies.....                                               | 7           |
| 2.1.1. WiFi Multimedia.....                                                  | 8           |
| 2.1.2. Summary.....                                                          | 8           |
| 2.2. Access Network – DOCSIS.....                                            | 8           |
| 2.2.1. OFDM.....                                                             | 8           |
| 2.2.2. MAP Intervals.....                                                    | 8           |
| 2.2.3. Upstream AQM .....                                                    | 9           |
| 2.2.4. Downstream AQM.....                                                   | 10          |
| 2.2.5. Extra UDP Traffic .....                                               | 10          |
| 2.2.6. Summary.....                                                          | 11          |
| 2.3. Core Network .....                                                      | 11          |
| 2.3.1. Route and Metric Tuning.....                                          | 11          |
| 2.3.1. Disabling ECMP For Peering and Transit Links (Latency Stability)..... | 12          |
| 2.3.2. Summary.....                                                          | 13          |
| 3. Upcoming Technologies And Methods For Improving Latency And Jitter .....  | 13          |
| 3.1. In-Home Technologies .....                                              | 13          |
| 3.2. Access Network – DOCSIS.....                                            | 14          |
| 3.2.1. Low Latency DOCSIS .....                                              | 14          |
| 3.3. Core Network .....                                                      | 16          |
| 4. Conclusion.....                                                           | 17          |
| Abbreviations .....                                                          | 18          |
| Bibliography & References.....                                               | 19          |

## List of Figures

| Title                                                                                      | Page Number |
|--------------------------------------------------------------------------------------------|-------------|
| Figure 1 – Sample ISP Gaming/Latency Website (TELUS, 2020) .....                           | 4           |
| Figure 2 – Speedtest.net’s Result Page Showing Latency Result (Ookla, 2020).....           | 5           |
| Figure 3 – League of Legends’ Lag Report (Riot Games, 2020).....                           | 5           |
| Figure 4 – League of Legends’ ISP Leaderboard for Calgary, AB (Riot Games, 2020).....      | 6           |
| Figure 5 – Average Latency And Jitter.....                                                 | 6           |
| Figure 6 – Where’s the Latency? .....                                                      | 7           |
| Figure 7 – TCP Latency When AQM Is Enabled Or Disabled .....                               | 9           |
| Figure 8 – Average Latency Under Congestion.....                                           | 10          |
| Figure 9 – Latency And Jitter In The Presence Of A Second UDP Stream.....                  | 11          |
| Figure 10 – Latency Of Multiple Paths Through A Network.....                               | 12          |
| Figure 11 – Network Exit Points.....                                                       | 13          |
| Figure 12 – Dual Queue Service Flows (White, Sundaresan, & Briscoe, 2019) .....            | 14          |
| Figure 13 – DOCSIS Grant Delay Cycle (White, Sundaresan, & Briscoe, 2019).....             | 15          |
| Figure 14 – MAP Interval and MAP Processing Time (White, Sundaresan, & Briscoe, 2019)..... | 16          |
| Figure 15 – Relative Latency Improvements .....                                            | 18          |

## List of Tables

| <b>Title</b>                                                                     | <b>Page Number</b> |
|----------------------------------------------------------------------------------|--------------------|
| Table 1 – Latency And Jitter In Relation To OFDM Channel Presence And Size ..... | 8                  |
| Table 2 – Latency And Jitter when MAP Interval Is Changed .....                  | 9                  |
| Table 3 – Latency When AQM Is Enabled/Disabled.....                              | 10                 |

## 1. Introduction


The gaming industry continues to grow, eclipsing the movie entertainment industry in revenues. In 2019, the digital video game market garnered \$120.1 billion (SuperData, 2020), while the worldwide box office earned \$42.5 billion (Comscore, 2020). In a 2018 survey, it was estimated that more than 23 million Canadians played video games (The Entertainment Software Association of Canada, 2018). This means more than 60% of Canadians could be considered gamers. Gamers are demanding higher speeds and lower latencies from their providers to achieve competitive and enjoyable multiplayer gaming experiences. These customers are choosing their service provider on either real - or perceived - advantages that the network technologies offered by providers. Additionally, real time applications such as voice and video conferencing require low latency and jitter to perform with high reliability and quality. While it is impossible for any single service provider to control end to end latency of internet traffic beyond their own network, there are steps that can be taken to achieve the lowest possible latency. Multiple Service Operators (MSOs) are constantly developing and deploying the latest technologies and strategies to improve latency and jitter on their networks.

ISPs advertise technologies that reduce latency, and third party measurements that indicate lower latency than competitors. For example, TELUS, a Canadian telco ISP, has a website dedicated to explaining latency and claiming that Fibre To The Home (FTTH) is better for gaming.



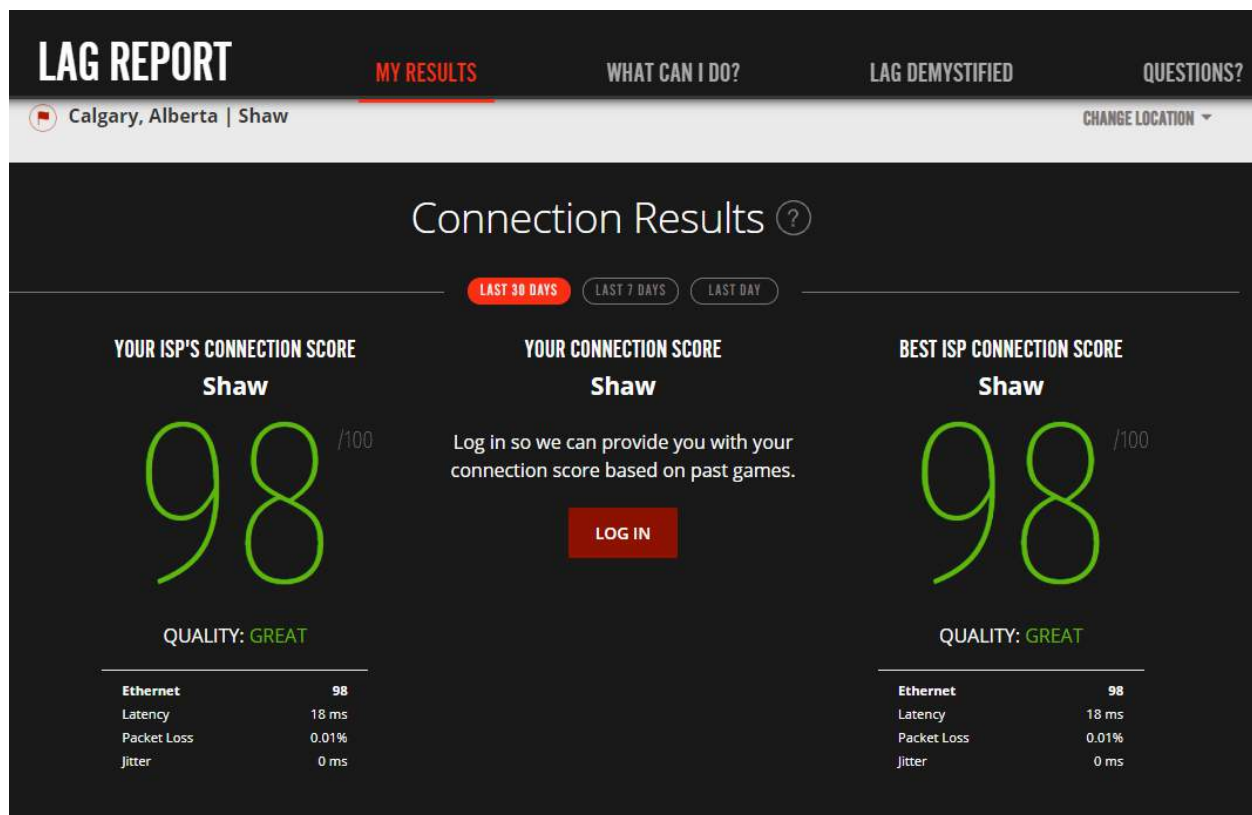
**Figure 1 – Sample ISP Gaming/Latency Website (TELUS, 2020)**

Speed test sites also are starting to show latency scores as another tool for measuring ISP quality (Speedtest.net / Ookla, DSLReports speedtest, and bing.com's speedtest to name a few).

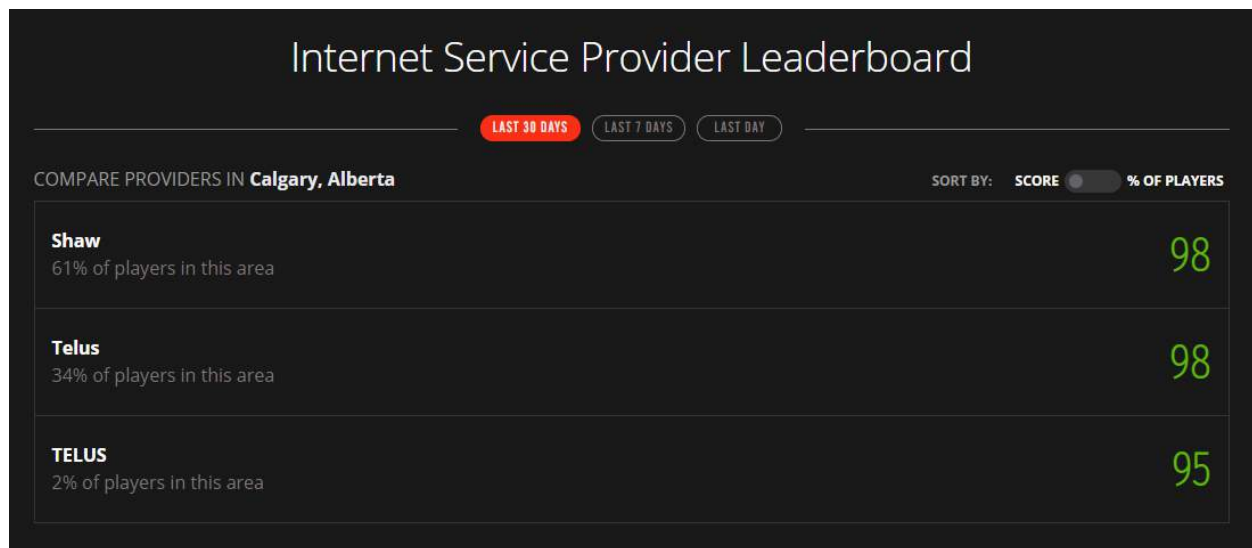


**Figure 2 – Speedtest.net’s Result Page Showing Latency Result (Ookla, 2020)**

League of Legends publishes ISP connectivity scores based on latency, jitter and packetloss. Their North American site for this is located at <https://lagreport.na.leagueoflegends.com/en/> (Figure 3 – League of Legends’ Lag Report) and shows not only your current provider's score, but a "leaderboard" of providers in your area (Figure 4 – League of Legends’ ISP Leaderboard for Calgary, AB), and many games have in-game displays of latency statistics.




**Figure 3 – League of Legends’ Lag Report (Riot Games, 2020)**



**Figure 4 – League of Legends’ ISP Leaderboard for Calgary, AB (Riot Games, 2020)**

Gamers are very interested in latency and jitter because of its impact on targeting and player location within the game – which can mean the difference between winning and losing. Jitter, which is often unpredictable, can be especially challenging for game servers to compensate for.

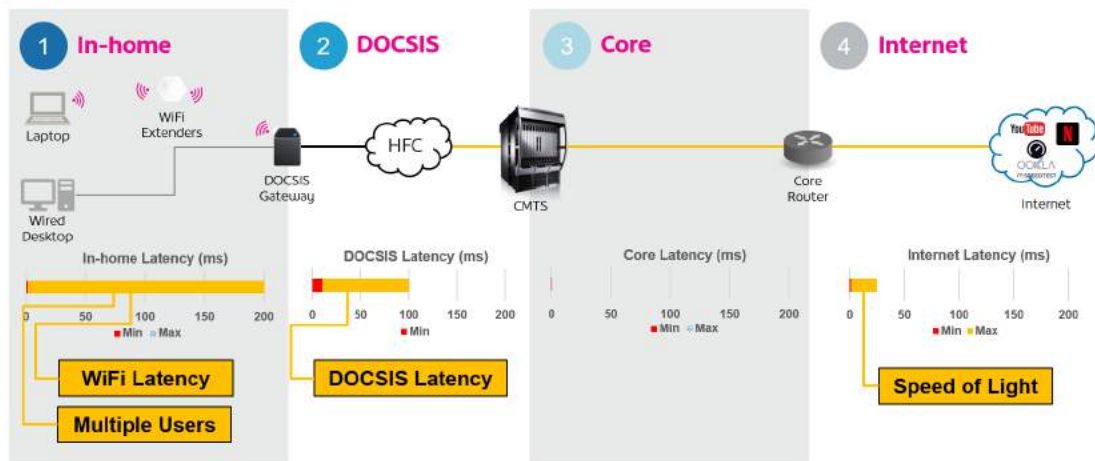
While average latency is the amount of time (on average) a packet takes to traverse the network, it is measureable and relatively stable (on average). Jitter is the difference in latency between different packets and is the variability of latency. The more variability, the harder it is for a game to be accurate and consistent between players.



**Figure 5 – Average Latency And Jitter**

This paper will discuss current and future methods that help improve latency and jitter at several points in the network: in the home, in the Data Over Cable Service Interface Specification (DOCSIS) network, and in the core ISP network. Technologies such as WiFi 6 Orthogonal Frequency-Division Multiple Access (OFDMA), CableLabs' Dual Channel WiFi, Low Latency DOCSIS (LLD) and routing optimizations will be discussed. We will show the level of improvement available at each part of the network, and the perceived and actual performance gains that can be expected, as well as the promise of LLD in the DOCSIS network.

## Where's The Latency



Shaw | Freedom

Figure 6 – Where's the Latency?

We will discuss three of the main areas shown in Figure 5; In-Home, the DOCSIS network, and the Core network, discussing some currently-available methods for reducing latency in these areas. Then, this paper will look at some technologies coming in the near future for each of these areas.

## 2. Current Methods For Improving Latency and/or Jitter

### 2.1. In-Home Technologies

Latency in the home is usually caused by WiFi usage. When wired, most home networks perform in the millisecond range, and quite frequently under a millisecond of latency is added traversing the home network. Over WiFi, conditions can vary wildly in different locations within the home, sometimes even in the same location. Yet customers continue to use WiFi as a primary method for connecting devices due to ease of use and the lack physical wiring in the home.



### 2.1.1. WiFi Multimedia

WiFi Multimedia (WMM) is a WiFi specification that was built to prioritize voice and video traffic over best effort internet traffic and also allows to de-prioritize background, non-latency sensitive traffic. When congestion happens WMM can ensure prioritized traffic will continue to get the bandwidth it requires. This also means it spends less time in the queue, avoiding latency and jitter penalties. We have seen applications such as Zoom employ WMM in the wild today, ensuring video calls perform well even in hostile WiFi environments. However, this is all at the cost of applications and traffic that don't use WMM, as their traffic will be penalized in WiFi congestion/contention scenarios. Any router supporting 802.11n or later should also have WMM enabled by default.

### 2.1.2. Summary

For applications that utilize WMM, and during congestion, the improvement can be tens to hundreds of milliseconds, as non-WMM WiFi queues can be quite large.

## 2.2. Access Network – DOCSIS

Technologies that can be employed in today's DOCSIS network include enabling Active Queue Management (AQM), adding orthogonal frequency-division multiplexing (OFDM) for DOCSIS 3.1 modems, and reducing DOCSIS Upstream Bandwidth Allocation Map (MAP) message intervals to 1ms or less. For the following tests, the test suite used defines jitter as IP Packet Delay Variation (IPDV) and is the average of all differences between each packet sent.

### 2.2.1. OFDM

Adding even a relatively small OFDM channel (64MHz) can allow for a small (0.5ms) end to end latency improvement. Based on Shaw's testing to date, this improvement doesn't seem to vary much with the size of the OFDM channel; it is only a function of whether the OFDM channel is present (and operational) or not.

**Table 1 – Latency And Jitter In Relation To OFDM Channel Presence And Size**

| OFDM size (MHz)    | Average Latency (ms) | Average Jitter (ms) |
|--------------------|----------------------|---------------------|
| 0 (not configured) | 9.62627              | 2.19518             |
| 64                 | 8.93366              | 2.22067             |
| 192                | 8.93708              | 2.21159             |

### 2.2.2. MAP Intervals

The MAP messages control when modems can request bandwidth, and also tells modems when they can transmit the data they have. Reducing the time between these messages can have an impact on latency as it gives more opportunities for modems to request bandwidth. Reducing the MAP interval to 1 ms or less is a requirement of LLD, but it is also something that is tunable today on some CCAP hardware.


Based on Shaw's testing to date, moving to < 1 ms MAP intervals can lower average latency a further 0.2ms, however it negatively impacts jitter by about about 0.2 ms as well.

**Table 2 – Latency And Jitter when MAP Interval Is Changed**

| MAP Interval (ms) | Average Latency (ms) | Average Jitter (ms) |
|-------------------|----------------------|---------------------|
| 3.2               | 10.06525             | 1.93527             |
| 0.8               | 9.86053              | 2.20513             |

### 2.2.3. Upstream AQM

Active Queue Management helps protect any network queue (wherever it is enabled) from ping spikes and latency due to network queuing. Commonly referred to as “buffer bloat”, this problem occurs when a bottleneck on the network receives data faster than it can send it out; it causes network latency to climb. Some of these queues can hold a large amount of data, allowing them to hold packets for very long times; sometimes into the range of multiple seconds. AQM uses the TCP mechanism of dropping packets to slow down TCP streams and actively monitors the time packets are spending in the queue. DOCSIS cable modems utilize the DOCSIS PIE algorithm to determine which packets are required to be dropped. Figure 7 shows the improvement with AQM enabled vs. disabled. As mentioned, since TCP dropping is the mechanism used to improve latency, when the upstream is congested with pure UDP traffic, the improvement is negligible. However, in a mixed or pure TCP traffic scenario, the latency can be improved dramatically, as seen in Table 3.




**Figure 7 – TCP Latency When AQM Is Enabled Or Disabled**

**Table 3 – Latency When AQM Is Enabled/Disabled**

|                                          | Commscope<br>TG3482 | Technicolor<br>CGM4140COM | Commscope<br>CM8200A | Hitron<br>CODA-46 | Hitron<br>CODA-4582<br>Bridged | Hitron<br>CODA-4582<br>Gateway | Technicolor<br>TC4400-AM |
|------------------------------------------|---------------------|---------------------------|----------------------|-------------------|--------------------------------|--------------------------------|--------------------------|
| <b>AQM Enabled<br/>UDP Traffic (ms)</b>  | 154.1               | 181.0                     | 174.6                | 168.3             | 421.1                          | 8.5                            | 172.9                    |
| <b>AQM Disabled<br/>UDP Traffic (ms)</b> | 201.8               | 216.2                     | 213.2                | 194.1             | 22212.7                        | 14.7                           | 209.0                    |
| <b>AQM Enabled<br/>TCP Traffic (ms)</b>  | 26.2                | 23.3                      | 24.0                 | 25.6              | 21.5                           | 21.3                           | 22.2                     |
| <b>AQM Disabled<br/>TCP Traffic (ms)</b> | 120.8               | 122.0                     | 126.1                | 196.6             | 124.8                          | 127.9                          | 129.9                    |

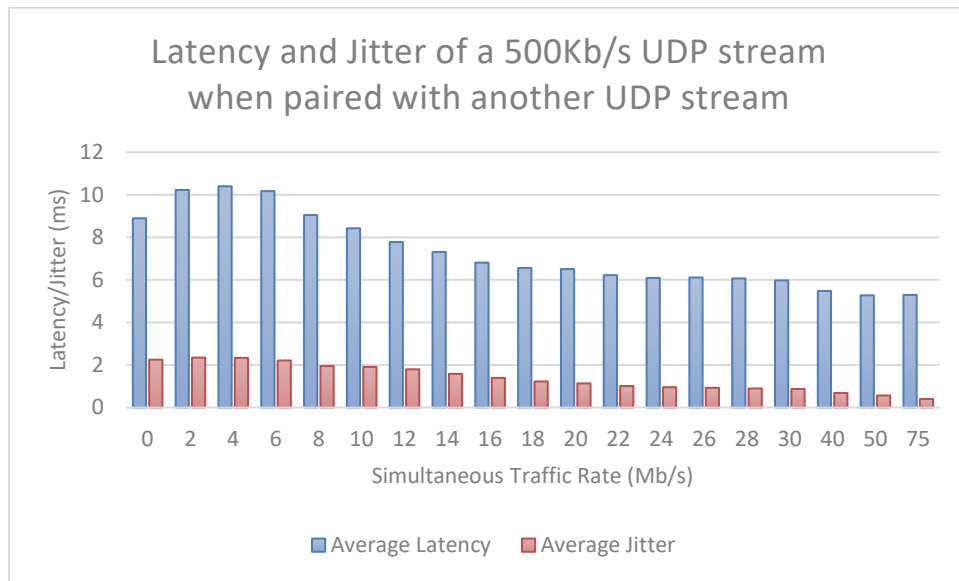
#### 2.2.4. Downstream AQM

One of our CCAP vendors implements the Weighted Random Early Detection (WRED) AQM algorithm for downstream traffic. However, by default, it is turned off, so when customers saturate their downstream service flows with downstream traffic, latency increases due to bufferbloat on the service flow. We configure our downstream to use the WRED AQM with a minimum threshold of 40 ms, a maximum threshold of 60 ms, and drop probability of 50%. Figure 8 shows the difference in latency over a short (about 1 minute) download between WRED being enabled and disabled.

**Figure 8 – Average Latency Under Congestion**

#### 2.2.5. Extra UDP Traffic

As more traffic bandwidth is injected through a set of service flows, latency and jitter appear to start to at first increase, and then decrease for a final improvement of over 3.5ms in latency, and over 1.6ms in jitter. This experiment added a second 75Mb/s UDP stream, and compared to the results to cases where no extra data streams were added.



**Figure 9 – Latency And Jitter In The Presence Of A Second UDP Stream**

This is clearly not practical for a production deployment, as it locks up a lot of upstream bandwidth. It may also just be an artifact of the scheduler used by the specific CCAP chassis used in this test, but is nonetheless interesting for its results in lowering both latency and improving jitter.


### **2.2.6. Summary**

Utilizing DOCSIS 3.1 channels (OFDM) and reducing map intervals to 1 ms or less can bring latency down by about 0.7 ms combined, at the cost of some jitter (increase of an average jitter of about 0.2 ms). Adding AQM can also ensure that latency does not grow unreasonably due to congestion, and it can also save over 100 ms of latency during congested periods.

## **2.3. Core Network**

### **2.3.1. Route and Metric Tuning**

This method can have a significant impact on latency. Results can have a greater than 10 ms improvement over non-optimal routes. Mapping is done by using ICMP to find the latencies of different routes between different paths through the network. An example is shown in Figure 10.



**Figure 10 – Latency Of Multiple Paths Through A Network**


The yellow path represents the network paths when no metric tuning has been done, and all links are equal cost. In this case, there are two paths. Depending on the load balancing algorithm the first router on the right uses, or could use one or the other path (39 or 40 ms), or could load balance packets between them for a single session resulting in a 1 ms jitter on the path.

When metrics are tuned to take the latency of each link into account, the resulting path is the blue path, giving the customer 36 ms of latency to traverse the network to the peer/transit point, and no jitter.

### **2.3.1. Disabling ECMP For Peering and Transit Links (Latency Stability)**

Border Gateway Protocol (BGP) will sometimes advertise equal costs for multiple peering or transit points for a given network. In this case, latencies can vary wildly from session to session as traffic moves between these different peering or transit links. End users would complain that latency was “good” on one day, but was “bad” on another day as their gaming session would change the exit point from our network.

In Figure 11, with Equal Cost Multipath (ECMP) enabled, a customer could reach a server through two network paths in the local network, if BGP had the same metric at each edge router. If the total Return Trip Time (RTT) for the yellow path is noticeably different than the green path, the customer will notice that for different gaming sessions to this server, he could see different latencies. Allowing the IGP metrics from metric tuning to be the tiebreaker for routing means that the network path (in this case) would always be the green path due to the lower latency internally. Depending on the latency to the server outside the network, the total RTT could actually be higher through the green path, but the latency being stable across sessions was more acceptable to customers, even if it was higher for specific servers.



**Figure 11 – Network Exit Points**

### **2.3.2. Summary**

These methods in the Core Network can improve latencies from 1 to 50 ms or more, depending on size and complexity of the network.

## **3. Upcoming Technologies And Methods For Improving Latency And Jitter**

### **3.1. In-Home Technologies**

WiFi 6 includes Orthogonal Frequency-Division Multiple Access (OFDMA). This feature improves bandwidth sharing with large numbers of clients to prevent clients talking over each other, thereby reducing latency and jitter in busier networks. In order to take advantage of OFDMA, however, all clients must be WiFi 6 enabled with OFDMA capability.

Dual Channel WiFi™ is a technology developed by CableLabs to provide a separate (or multiple separate) downstream channels in combination with a “legacy” channel. This allows large downloads, and streaming to be shunted to this separate channel, keeping the legacy channel open for upstream and legacy clients, reducing latency and jitter by reducing contention and congestion when clients are requesting large amounts of traffic.

## 3.2. Access Network – DOCSIS

### 3.2.1. Low Latency DOCSIS

The two technologies in LLD that will contribute the most to latency and jitter improvements are dual-queue and the Proactive Grant Service (PGS). LLD AQM will also ensure that latency is held down during congestion and Coupled AQM will ensure that bandwidth is available to both the low latency queue as well as the classic queue during congestion

#### 3.2.1.1. Dual-Queue

Dual-queue provides a separate set of DOCSIS service flows, so time sensitive traffic doesn't have to compete with non time sensitive traffic within a single queue. This means that even when there is a large amount of traffic on the legacy service flow, it will not affect the queuing behaviour of the low latency service flow.

Dual-queue also allows the DOCSIS scheduler to schedule latency sensitive traffic differently to provide low latency and low jitter.

So, why don't we have all traffic low latency? Queuing delays are actually caused by some applications sending traffic in a manner that results in a build up of packets in the network. The biggest source of this type of traffic is any application that uses TCP today. Those applications need the network to provide a deep buffer to absorb bursts of traffic, and don't perform well if the network has shallow buffers. But games and latency sensitive applications typically deal with small packet loss better than latency variation, and large buffers allow for large variation or spikes in latency, as well as large latency if the buffers are kept full for a long time. Having small buffers just for latency sensitive applications will ensure packets are not held for too long. The dual-queue feature of Low Latency DOCSIS equipment gives both types of traffic the appropriate buffer for their needs: a shallow buffer for Low Latency traffic, and a deep buffer for Classic traffic.





Figure 12 – Dual Queue Service Flows (White, Sundaresan, & Briscoe, 2019)

This does mean that we will need to differentiate the traffic via some means. We do have access to DOCSIS packet classifiers, but these would be hard to manage as it would require classification of every type of traffic and every different game to be put into the low-latency service flow. LLD will use several

methods to classify traffic into the low latency service flow; a Differentiated Services (DiffServ) value or Explicit Congestion Notification (ECN). CableLabs has proposed a DiffServ value of 0x2A be defined as Non-Queue-Building (NQB), and LLD would use this value, as well as ECN values. As ECN is supported, Low-Latency Low-Loss Scalable throughput (L4S) will also be supported for applications that need both high bandwidth and low latency. CableLabs has been working with game developers to inform them how to mark their packets through gaming conferences and through the website pingspikeskill.com.

### 3.2.1.2. Proactive Grant Service

Proactive Grant Service (PGS) is a new DOCSIS upstream scheduling service included in LLD. PGS is like the Unsolicited Grant Service (UGS) used primarily for voice services, but where UGS is very static with the size of packets and the number of grants given, PGS is much more dynamic, allowing for different size packets and different amounts of bandwidth. PGS helps reduce the grant-delay cycle inherent in the DOCSIS protocol (Figure 13 – DOCSIS Grant Delay Cycle).



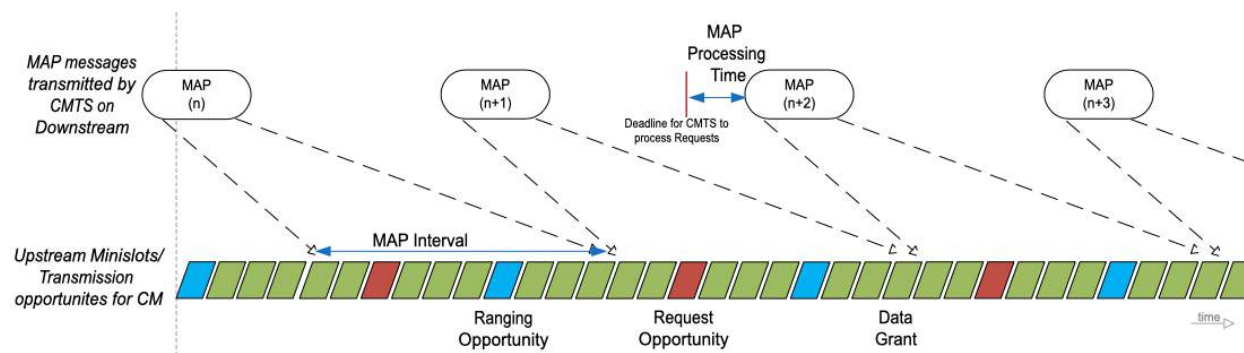
**Figure 13 – DOCSIS Grant Delay Cycle (White, Sundaresan, & Briscoe, 2019)**

As the name indicates, PGS proactively tries to predict and give grants for the modem to transmit data based on past behaviour, so that the modem doesn't have to request the grants (effectively removing the grant delay cycle latency); but the downside is that if the modem doesn't have any data to transmit, that bandwidth is now wasted.

### 3.2.1.3. Upstream Scheduling Improvements

Typically it takes more time than a single MAP interval to process a request for bandwidth, so reducing the MAP interval alone does not provide a huge boost to latency reduction. If the turnaround for processing the request can be reduced along with the MAP interval reduction, this should allow for a good reduction of latency and jitter even without PGS.





**Figure 14 – MAP Interval and MAP Processing Time (White, Sundaresan, & Briscoe, 2019)**

#### **3.2.1.4. Coupled AQM**

Coupled AQM ensures that there is shared bandwidth for both the low latency service flow and the classic service flow within the aggregate service flow during times of queuing. The total bandwidth consumed by the coupled service flows is limited by the AQM configuration settings. This feature ensures that the low latency service flow does not starve out the classic service flow under this condition, and aims to ensure that classic TCP sessions and future L4S TCP sessions all receive a fair bandwidth allocation.

#### **3.2.1.5. Summary**

Total potential improvements with LLD will run in the range of 5-10 ms of latency (more in cases with congestion/queuing), as LLD can deliver DOCSIS latency of ~1ms (White, Sundaresan, & Briscoe, 2019) and jitter will most certainly be improved as well with both PGS and upstream scheduling improvements. This will significantly close the “latency gap” that last mile fibre products use in their marketing, even though 5-10 ms may seem like a small improvement overall.

### **3.3. Core Network**

Currently the optimization of routing configurations for specific destinations is very manual and labour intensive, and it has only been done a handful of times when teams have the cycles to perform the required tuning. Going forward we are considering launching a project to keep these up to date and ensure they don’t cause problems as our network evolves and changes. We are also looking at enabling partial automation for these optimisations, but that is further in the future.

CIN networks can add another level of complexity to these calculations especially for more remote sites due to the diverse path requirements and distances involved. Even within a single metropolitan area, fiber paths and switching delays can contribute to a slight difference in latency between paths. Paying attention to the different latencies between each path will continue to be an ongoing process.

Cloud Gaming, Augmented Reality (AR) and Virtual Reality (VR) applications will start to require lower latency, and since the networks between users and where servers are hosted can make up a large portion of the latency budget, edge computing or edge cloud, where servers are moved into ISP facilities, or very close to them, can also play a part in the ability for an ISP to lower latency for their customers.


## 4. Conclusion

Carefully focusing on the parts of the network over which you have control provides ample opportunity for reducing inefficiencies. Enabling AQM and having devices and applications that utilize WMM can yield many latency benefits. In addition, wiring as many devices as possible in the home can help keep latency down in times of congestion. Enabling DOCSIS 3.1 downstream channels (OFDM) and reducing upstream MAP intervals to 1 millisecond or less can bring DOCSIS access latency down a small but measurable amount. Finally, ensuring your core network is optimized for latency via continual measuring and adjusting of IGP metrics can give you an edge over even FTTH.

Some technologies look impressive, but need to be framed within their use case. For example, AQM, WMM, WiFi6 OFDMA, and Dual Channel WiFi all perform impressively during congestion or on busy networks, but during normal operations may not yield much improvement for latency or jitter. That doesn't mean they aren't important, as even home networks can have short bursts of extreme traffic or a large number of devices requiring network resources at the same time. Some of the technologies may seem like they don't contribute much to a lower RTT on the network, such as 1 ms MAP Intervals, OFDM, or even LLD, but these can lower the difference of latencies seen between FTTH and DOCSIS, and in the case of LLD even eliminate that difference altogether.

In the future WiFi 6 within the home will further hold latency down as more and more devices are added to home networks. LLD may bring latency down to as low as 1 ms in the access network. Further optimizations on the Core network can ensure latency on a DOCSIS network is as low as even the best FTTH network.

Figure 15 below shows the rough amount of latency that can be shaved off the RTT with each technology or optimization. Some are additive, some (such as LDD, which includes MAP Intervals and AQM) includes other technologies listed. Some may seem small, but work under most conditions (MAP Intervals/OFDM) while others (marked with asterisks) seem like large gains, but only work under some conditions (AQM). LLD has technologies that will improve both areas, as seen with its two entries. These will all work together to improve latency and jitter on the network keeping latency sensitive applications and gamers satisfied and happy.



**Figure 15 – Relative Latency Improvements**

## Abbreviations

|        |                                                 |
|--------|-------------------------------------------------|
| AQM    | Active Queue Management                         |
| AR     | Augmented Reality                               |
| BGP    | Border Gateway Protocol                         |
| DOCSIS | Data Over Cable Service Interface Specification |
| FTTH   | Fiber To The Home                               |
| IGP    | Interior Gateway Protocol                       |
| ISP    | Internet Service Provider                       |
| LLD    | Low Latency DOCSIS                              |
| MAP    | Upstream Bandwidth Allocation Map               |
| OFDM   | Orthogonal Frequency-Division Multiplexing      |
| OFDMA  | Orthogonal Frequency-Division Multiple Access   |
| PGS    | Proactive Grant Service                         |
| RTT    | Round Trip Time                                 |
| WRED   | Weighted Random Early Detection                 |
| UGS    | Unsolicited Grant Service                       |
| VR     | Virtual Reality                                 |
| WMM    | WiFi Multimedia                                 |

## Bibliography & References

- Comscore. (2020, 01 10). *Comscore Reports Highest Ever Worldwide Box Office*. Retrieved from Comscore: <https://www.comscore.com/Insights/Press-Releases/2020/1/Comscore-Reports-Highest-Ever-Worldwide-Box-Office>
- Ookla. (2020, 07 27). *Speedtest by Ookla results*. Retrieved from Speedtest by Ookla : <https://www.speedtest.net/result/9820878391>
- Riot Games. (2020, 08). *League of Legends Lag Report*. Retrieved from League of Legends: <https://lagreport.na.leagueoflegends.com/en/>
- SuperData. (2020). *2019 Year In Review*. Retrieved from SuperData Research: <https://www.superdataresearch.com/2019-year-in-review/>
- TELUS. (2020, 08 01). *Best Internet for Gaming - Fibre Internet | TELUS*. Retrieved from TELUS: <https://www.telus.com/en/internet/best-internet-for-gaming>
- The Entertainment Software Association of Canada. (2018). *Essential Facts About The Canadian Video Game Industry*. Retrieved from The Entertainment Software Association of Canada: [http://theesa.ca/wp-content/uploads/2018/10/ESAC18\\_BookletEN.pdf](http://theesa.ca/wp-content/uploads/2018/10/ESAC18_BookletEN.pdf)
- White, G., Sundaresan, K., & Briscoe, B. (2019). *Low Latency DOCSIS: Overview And Performance Characteristics*. SCTE.

# **Distributed Gain Architecture**

## **Increased Performance, Decreased Power Draw**

A Technical Paper prepared for SCTE•ISBE by

**Nader Foroughi**

Senior Network Architect  
Shaw Communications  
2728 Hopewell Place NE, T1Y 7J7  
+1 403 648 5937  
[nader.foroughi@sjrb.ca](mailto:nader.foroughi@sjrb.ca)

**Jason Rupe**

Principal Architect  
CableLabs®  
858 Coal Creek Circle, Louisville, CO 80027  
303.661.3332  
[j.rupe@cablelabs.com](mailto:j.rupe@cablelabs.com)

# Table of Contents

| Title                                                                           | Page Number |
|---------------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                            | 5           |
| 2. Technological and Operational Challenges with Extended Spectrum DOCSIS ..... | 6           |
| 2.1. Plant Spacing and Drop-In Upgrades .....                                   | 6           |
| 2.2. Total Composite Power (TCP) .....                                          | 6           |
| 2.3. Taps .....                                                                 | 7           |
| 3. Cascaded Plant Design Challenges.....                                        | 8           |
| 3.1. Noise .....                                                                | 8           |
| 3.2. Distortion .....                                                           | 9           |
| 3.3. Designing a Noise-Limited System .....                                     | 9           |
| 4. Plant Models.....                                                            | 10          |
| 4.1. Traditional Node and Amplifier Outputs .....                               | 13          |
| 4.2. Traditional Node and Amplifier Noise Figure (NF) .....                     | 14          |
| 4.3. Modem (MDM) Transmit Power: .....                                          | 14          |
| 4.4. Modulation Order vs. Power and C/N .....                                   | 15          |
| 5. DGA Amplifier Considerations .....                                           | 16          |
| 5.1. DS Gain .....                                                              | 16          |
| 5.2. US Gain .....                                                              | 17          |
| 6. Traditional Plant DS Results .....                                           | 18          |
| 6.1. 135' Plant MDM Rx Powers .....                                             | 21          |
| 6.2. 190' Plant MDM Rx Powers .....                                             | 22          |
| 6.3. 204' Plant MDM Rx Powers .....                                             | 23          |
| 7. Booster Amplification DS Results .....                                       | 24          |
| 7.1. 190' and 204' Plant Design with Booster Amplifiers .....                   | 25          |
| 190' Plant Traditional Amplifier Rx Power with Booster Amplification: .....     | 26          |
| 204' Plant Traditional Amplifier Rx Power with Booster Amplification: .....     | 26          |
| 7.2. Booster Amplification Observations .....                                   | 28          |
| 8. DGA Design and DS Results .....                                              | 28          |
| 8.1. 135' Plant Design and DS Results .....                                     | 29          |
| 8.2. 190' Plant Design and DS Results .....                                     | 32          |
| 8.3. 204' Plant Design and DS Results .....                                     | 34          |
| 8.4. DGA DS Design Observations .....                                           | 37          |
| 9. Upstream Analysis and Considerations .....                                   | 37          |
| 9.1. Higher MDM Transmit Levels in Lower US Splits .....                        | 39          |
| 9.2. Higher Return Gain in DGA Amplifiers:.....                                 | 41          |
| 10. Power Draw .....                                                            | 42          |
| 11. Plant Reliability .....                                                     | 43          |
| 12. Future Considerations .....                                                 | 46          |
| 12.1. Unity Gain, Distortions and Cascade Limits .....                          | 47          |
| 12.2. US Gain and Performance .....                                             | 48          |
| 12.3. Design Standards .....                                                    | 48          |
| 12.4. 3 GHz .....                                                               | 49          |
| 13. Conclusion.....                                                             | 49          |
| Abbreviations .....                                                             | 50          |
| Bibliography & References.....                                                  | 50          |

## List of Figures

| Title                                                                                              | Page Number |
|----------------------------------------------------------------------------------------------------|-------------|
| Figure 1 – Upstream Data Trends Pre and Post COVID-19 .....                                        | 5           |
| Figure 2 – TCP vs. MER .....                                                                       | 7           |
| Figure 3 – Signal Level Balanced Between Noise and Distortion.....                                 | 10          |
| Figure 4 – Plant Models .....                                                                      | 11          |
| Figure 5 – Analyzed Plant Types .....                                                              | 12          |
| Figure 6 – Node and Amplifier Output PSD .....                                                     | 13          |
| Figure 7 – MDM Output Power PSD .....                                                              | 15          |
| Figure 8 – DGA Amplifier DS Gain .....                                                             | 16          |
| Figure 9 – DGA Amplifier US Gain .....                                                             | 17          |
| Figure 10 – 135' Amplifier Rx Power @ Amp. Port .....                                              | 18          |
| Figure 11 – 190' Amplifier Rx Power @ Amp. Port .....                                              | 19          |
| Figure 12 – 204' Amplifier Rx Power @ Amp. Port .....                                              | 20          |
| Figure 13 – 135' Span 1 MDM Rx Power/6 MHz .....                                                   | 21          |
| Figure 14 – 135' Span 2 MDM Rx Power/6 MHz .....                                                   | 21          |
| Figure 15 – 190' Span 1 MDM Rx Power/6 MHz .....                                                   | 22          |
| Figure 16 – 190' Span 2 MDM Rx Power/6 MHz .....                                                   | 22          |
| Figure 17 – 204' Span 1 MDM Rx Power/6 MHz .....                                                   | 23          |
| Figure 18 – 204' Span 2 MDM Rx Power/6 MHz .....                                                   | 24          |
| Figure 19 – 190' and 204' Plant with Mid-Span Booster Amplification .....                          | 25          |
| Figure 20 – 190' Amplifier Rx Power @ Port with Booster Amplification.....                         | 26          |
| Figure 21 – 204' Amplifier Rx Power @ Port with Booster Amplification.....                         | 26          |
| Figure 22 – 190' Span 1&2 MDM Rx Power/6MHz .....                                                  | 27          |
| Figure 23 – 204' Span 1&2 MDM Rx Power/6MHz .....                                                  | 28          |
| Figure 24 – 135' DGA Conversion .....                                                              | 29          |
| Figure 25 – 135' Plant DGA Amplifiers' Rx Power @ Ports.....                                       | 30          |
| Figure 26 – DGA 135' Span 1 MDM Rx Power/6MHz .....                                                | 31          |
| Figure 27 – DGA 135' Span 2 MDM Rx Power/6MHz .....                                                | 31          |
| Figure 28 – 190' DGA Conversion .....                                                              | 32          |
| Figure 29 – 190' Plant DGA Amplifiers' Rx Power @ Ports.....                                       | 32          |
| Figure 30 – DGA 190' Span 1 MDM Rx Power/6 MHz .....                                               | 33          |
| Figure 31 – DGA 190' Span 2 MDM Rx Power/6 MHz .....                                               | 34          |
| Figure 32 – 204' DGA Conversion .....                                                              | 34          |
| Figure 33 – 204' Plant DGA Amplifiers' Rx Power @ Ports.....                                       | 35          |
| Figure 34 – DGA 204' Span 1 MDM Rx Power/6 MHz .....                                               | 36          |
| Figure 35 – DGA 204' Span 2 MDM Rx Power/6 MHz .....                                               | 36          |
| Figure 36 – 204' Plant Return Path DGA Amplifiers' Rx Power @ Ports .....                          | 37          |
| Figure 37 – 204' Plant Return Path Node Rx Power @ Port.....                                       | 38          |
| Figure 38 – Last Span DGA-to-Node Level Comparison – 2dB Insertion Loss Taps vs. Regular Taps .... | 39          |
| Figure 39 – Raised MDM Output Power PSD in Various Splits.....                                     | 40          |
| Figure 40 – Last Span DGA-to-Node Level Comparison – Raised MDM PSD .....                          | 40          |
| Figure 41 – High vs. Low Gain DGA Return Amplifier.....                                            | 41          |

|                                                                                                       |    |
|-------------------------------------------------------------------------------------------------------|----|
| Figure 42 – Last Span DGA-to-Node Level Comparison – High Gain vs. Low Gain Return DGA Amplifier..... | 42 |
| Figure 43 – DGA Availability .....                                                                    | 45 |
| Figure 44 – DGA Un-availability.....                                                                  | 45 |
| Figure 45 – Relationship between Cascade, Noise and Distortion.....                                   | 48 |


## List of Tables

| <b>Title</b>                                                           | <b>Page Number</b> |
|------------------------------------------------------------------------|--------------------|
| Table 1 – Node and Amplifier NF .....                                  | 14                 |
| Table 2 – MDM Tx Power/1.6 MHz.....                                    | 14                 |
| Table 3 – MDM Tx Power/6.4 MHz.....                                    | 14                 |
| Table 4 – MDM Tx Power/6.4 MHz.....                                    | 15                 |
| Table 5 – DGA Amplifier DS and US NF .....                             | 16                 |
| Table 6 – 135’ Traditional Amplifier Contributions to System CNR ..... | 18                 |
| Table 7 – 190’ Traditional Amplifier Contributions to System CNR ..... | 19                 |
| Table 8 – 204’ Traditional Amplifier Contributions to System CNR ..... | 20                 |
| Table 9 – Amplifier Contributions to System C/N .....                  | 27                 |
| Table 10 – 135’ DGA Amplifier Contributions to System C/N.....         | 30                 |
| Table 11 – 190’ DGA Amplifier Contributions to System C/N.....         | 33                 |
| Table 12 – 204’ DGA Amplifier Contributions to System C/N.....         | 35                 |
| Table 13 – Power Draw Comparisons .....                                | 43                 |



## 1. Introduction

As data trends and usage increase, operators are looking for methods to increase the capacity of the network. This is especially the case with upstream. During the COVID-19 pandemic, the need to increase the return-band spectrum bandwidth and throughput became evident. The graph below demonstrates this increase in usage:



**Figure 1 – Upstream Data Trends Pre and Post COVID-19**

As a part of upgrading the outside plant (OSP) to 1.8 GHz extended spectrum DOCSIS (ESD) and beyond, to achieve 10 Gbps and more, many operators face the costly prospect of amplifier re-spacing. Given how costly and labour intensive plant re-spacing can be, innovating ideas to overcome this challenge are highly encouraged.

MSOs have traditionally relied on high gain amplifiers to overcome coaxial loss in the access architecture. Traditional amplifiers have served this purpose well, providing 50-60 dBmV of gain, however, they can be power hungry, drawing 60-80 watts each. They can also introduce unwanted distortions in the spectrum, decreasing the signal quality and essentially lowering the achievable throughput in the network.

Distributed gain architecture involves deploying smaller and lower gain amplifiers in selective areas of the network, in conjunction or instead of high gain amplifiers. These amplifiers have a much lower power draw in comparison traditional amplifiers, which can drastically decrease the draw from the existing power supplies, leaving more room for other technologies to be deployed in the access network.

Due to the simplicity of these amplifiers, being single stage with a fixed gain and tilt, they can also be deployed in conjunction with traditional amplifiers. This can boost the end-of-line performance in areas

that the span loss of the plant cannot be overcome with the available total composite power (TCP) of the traditional amplifiers.

This paper will provide an elaborate study and comparison between a traditional N+2 plant, N+2 plant with booster amplification, and a fully distributed gain versions of the same plant models. An end-of-line performance and power analysis will be provided for each scenario.

## **2. Technological and Operational Challenges with Extended Spectrum DOCSIS**

Prior to the analysis, we must first discuss the challenges that operators face when considering upgrading their access networks to ESD.

In ‘traditional’ plant, being 750 MHz or even 1 GHz, most MSOs expect 4kQAM to be achievable by each orthogonal frequency division multiplexing (OFDM) carrier deployed, however, the same is not true for 1.8 GHz and beyond. One of the primary reasons for this would be the current spacing that the outside plant is designed to.


### **2.1. Plant Spacing and Drop-In Upgrades**

Most OSP architectures are designed to 550 MHz and ‘stretched’ to 1 GHz. As a result, most of the radio frequency (RF) power in plant actives, including nodes and amplifiers, has been utilized to overcome the existing span losses. Span loss is defined as the total insertion loss of all the elements in a hybrid-fibre-coax (HFC) span, measured in dB. This includes all the plant passives such as taps, splitters and couplers. Although the span losses today are manageable with the current amplifier gains, they will certainly become a major point of concern when the spectrum is expanded to higher frequencies. As an example, a span loss of 35 dB at 1 GHz in a traditional plant equates to 49 dB at 1.8 GHz.

Plant re-spacing is always an option, however it can be extremely costly and, as a result, operators will rely on the expanded power of amplifier gain chips to overcome span losses.

### **2.2. Total Composite Power (TCP)**

It is generally understood that 1.8 GHz amplifier chips will have ~75 dBmV of total composite power (TCP) available. With that in mind, not all of this power is available for use. As an example, the figure below demonstrates the trade-off between modulation error ratio (MER) and TCP utilized:



**Figure 2 – TCP vs. MER**

As a general rule of thumb, 3 dB of back-off is needed to achieve 40+ dB MER, which is typically what operators aim for. Along with that, the internal loss of the active device has to be accounted for, which is usually 2 dB. To summarize, there is a total of 70 dBmV of power to be utilized at the port of each active device. This can be a concern for operators given that 65-68 dBmV of TCP has already been allocated to overcome the span losses in the ‘traditional’ plant.

### **2.3. Taps**

Taps and passives can be another point of concerns when upgrading the OSP to 1.8 GHz. Traditionally, most operators have relied on face-plate upgrades to expand the spectrum range of plant taps and passives. This is generally accepted as a faster and more cost-effective method to upgrade the available bandwidth of taps.

Unfortunately, this might not be the case with 1.8 GHz upgrades. A face-plate upgrade of the current 1 GHz taps can potentially expand the bandwidth to approximately 1.6 GHz. It should also be noted that this is a best effort.

This can be a concern given the uncertainty of the maximum available bandwidth in the plant. As a result, it is generally accepted that taps and passives have to be swapped out for 1.8 GHz version. Given that the entire housing of the tap has to be swapped out as a part of this effort, most of the taps being developed will have housings that can support up to 3 GHz with future face-plate upgrades, future proofing the plant for 3 GHz upgrades.

### 3. Cascaded Plant Design Challenges

HFC architectures can be divided into two categories:

- Passive plant (N+0): where no amplifiers are used after the node
- Cascaded plant (N+X): where amplifiers are used to boost the signal multiple times to the end-of-line

When designing an N+0 plant, the main point of concern is the output performance of the node. Assuming that we are operating in a distributed access architecture (DAA) plant, the primary drivers for the plant quality would be the MER of the DAA device. Since no amplifiers are used to boost the signal, no noise or distortion is added to the primary signal being generated by the DAA device.

On the contrary, when designing an N+X plant, the following can be of concern:

- Amplifier noise contribution
- Amplifier distortion contribution

**Note:** In order to calculate the overall system carrier to noise ratio (C/N) a starting C/N has to be assumed. Due to the continued development in this area, no starting C/N has been assumed from the RF source (RPD or RMD). Instead, the cascaded amplifier network's contribution to the system C/N has been calculated in this paper. Once a starting C/N is determined at the output of the node, the overall system C/N can be calculated.

#### 3.1. Noise

Designing a cascaded system for optimal carrier to noise is always a big priority for an operator. One of the biggest contributors in system design is the receive power (Rx Power) at the amplifier, given that it is one of the primary drivers for the overall system C/N.

The equation below calculates the C/N of a single amplifier:

$$C/N (dB) = C_i(dBmV) + 57.4 - NF(dB)$$

Where:

- $C_i$ : input signal
- $NF$ : Noise figure of the amplifier

**Note:** the number 57.4 is the thermal noise power in dBmV for 6 MHz QAM carriers.

The equation above shows the significance of the Rx power versus noise figure of the amplifier, in overall system design.

The overall system C/N for amplifiers operating at different output levels can be derived from the following equation:

$$C/N_{total} (dB) = -10 \log \left\{ 10^{\frac{-C/N_1}{10}} + 10^{\frac{-C/N_2}{10}} + \dots + 10^{\frac{-C/N_n}{10}} \right\}$$

Where,  $C/N_x$  is the carrier to noise of each amplifier calculated independently.

When cascading identical amplifiers operating at the same output level, the following approximation is typically used:

$$C/N_{total} (dB) = C/N_x - 10\log n$$

Where:

- $C/N_x$ : the carrier to noise of a single amplifier
- $n$ : the number of identical amplifiers in cascade.

### 3.2. Distortion

The buildup of distortions in a cascaded plant is less predictable than noise. Knowing that almost all carriers deployed in the spectrum in the future will be digital, the distortion products can be summed into carrier to intermodulation noise (CIN), which will increase the noise level that should be considered in the system C/N. Due to lack of availability of data in this realm, CIN was not considered in this paper but it is something that needs to be studied extensively, discussed in section 12.

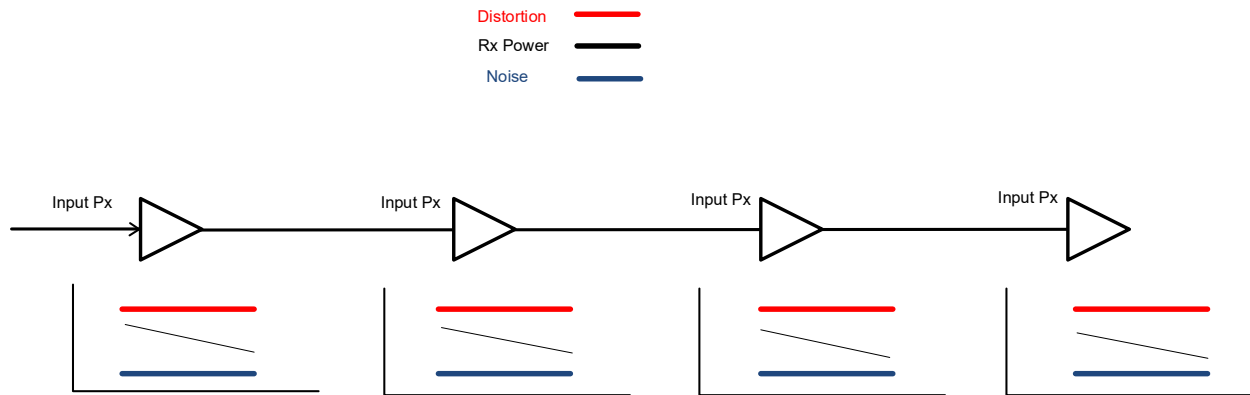
Since distributed gain amplifiers have very low distortion characteristics, due to the low gain and simplicity of these amplifiers, and the fact that distortions for traditional amplifiers are highly unpredictable, composite second order distortion (CSO), composite triple beat (CTB) and subsequently CIN have been nullified in the calculations for end-of-line performance.

### 3.3. Designing a Noise-Limited System

For optimal performance, operators design systems that are unity gain. This means that the loss between two amplifiers is equal to the gain of each amplifier. If the loss is less than the gain, output power needs to be increased and as a result, distortions will accumulate. In contrast, if the loss is greater than the gain, then the input power will be too low to the input of the amplifier, degrading the C/N of the system.

Due to the difficulties that come with designing a system that is both noise and distortion limited, removing one of those parameters will be optimal. Given that noise performance of amplifiers is far more straight-forward in comparison to distortion, designing a noise-limited system is an attractive idea.

Since distortions are highly dependent on output power TCP, designing a noise-limited system can be achieved by reducing the output power out of the node/amplifiers and making sure the signal is received at the next amplifier at a high enough level for acceptable C/N in spite of the amplifier's noise figure (NF).



**Figure 3 – Signal Level Balanced Between Noise and Distortion**

## 4. Plant Models

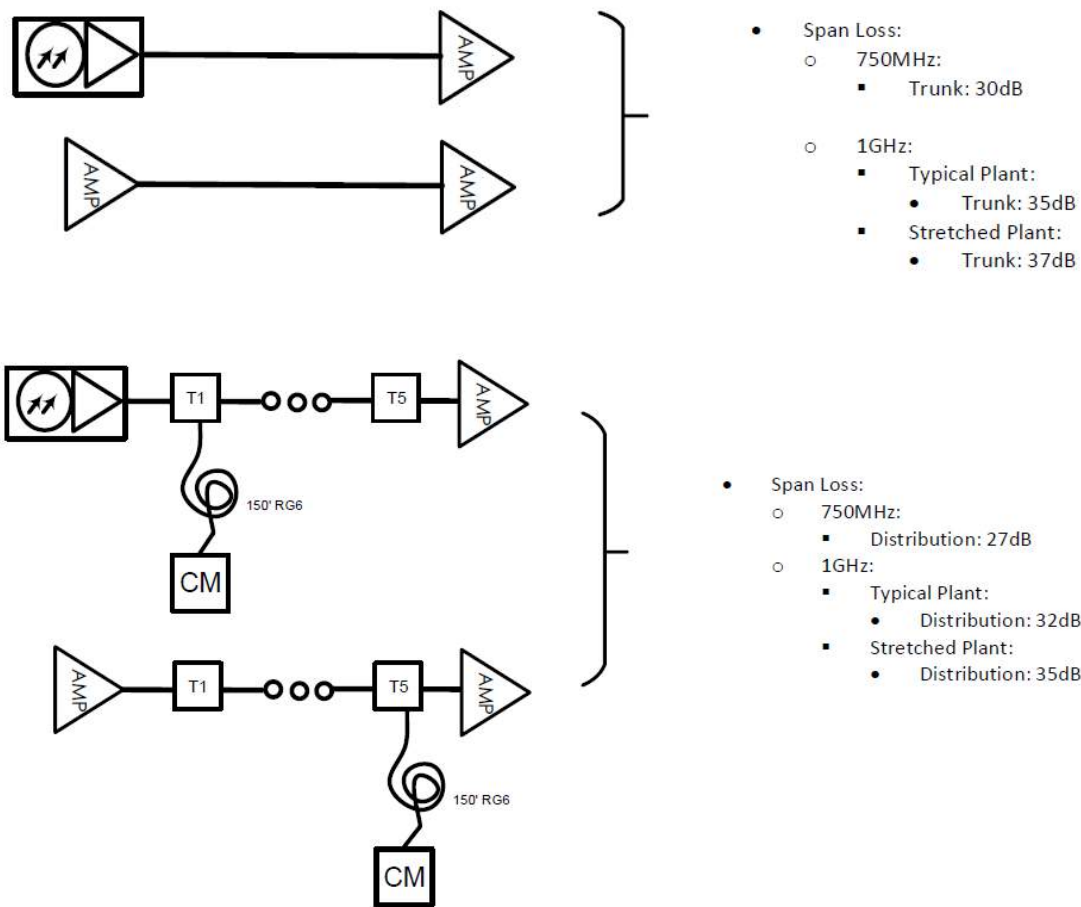
In order to encompass most of the HFC architectures deployed, the below plant models and assumptions were considered for this analysis.

Note: Trunk spans are defined as spans that are untapped. Distribution spans are tapped. Both trunk and distribution span losses include all other passive elements' insertion losses, such as splitters and couplers.

### Assumptions:

- Modem:
  - Point of entry (PoE) device
- Drop:
  - Cable: RG6
  - Length: 150 feet
- Number of taps in each span:
  - 5
- Distribution span losses at 1GHz:
  - Typical plant: 35 dB
  - Stretched plant: 37 dB
- Trunk span losses at 1 GHz:
  - Typical plant: 32 dB
  - Stretched plant: 35 dB

The figure below summarizes the parameters above:



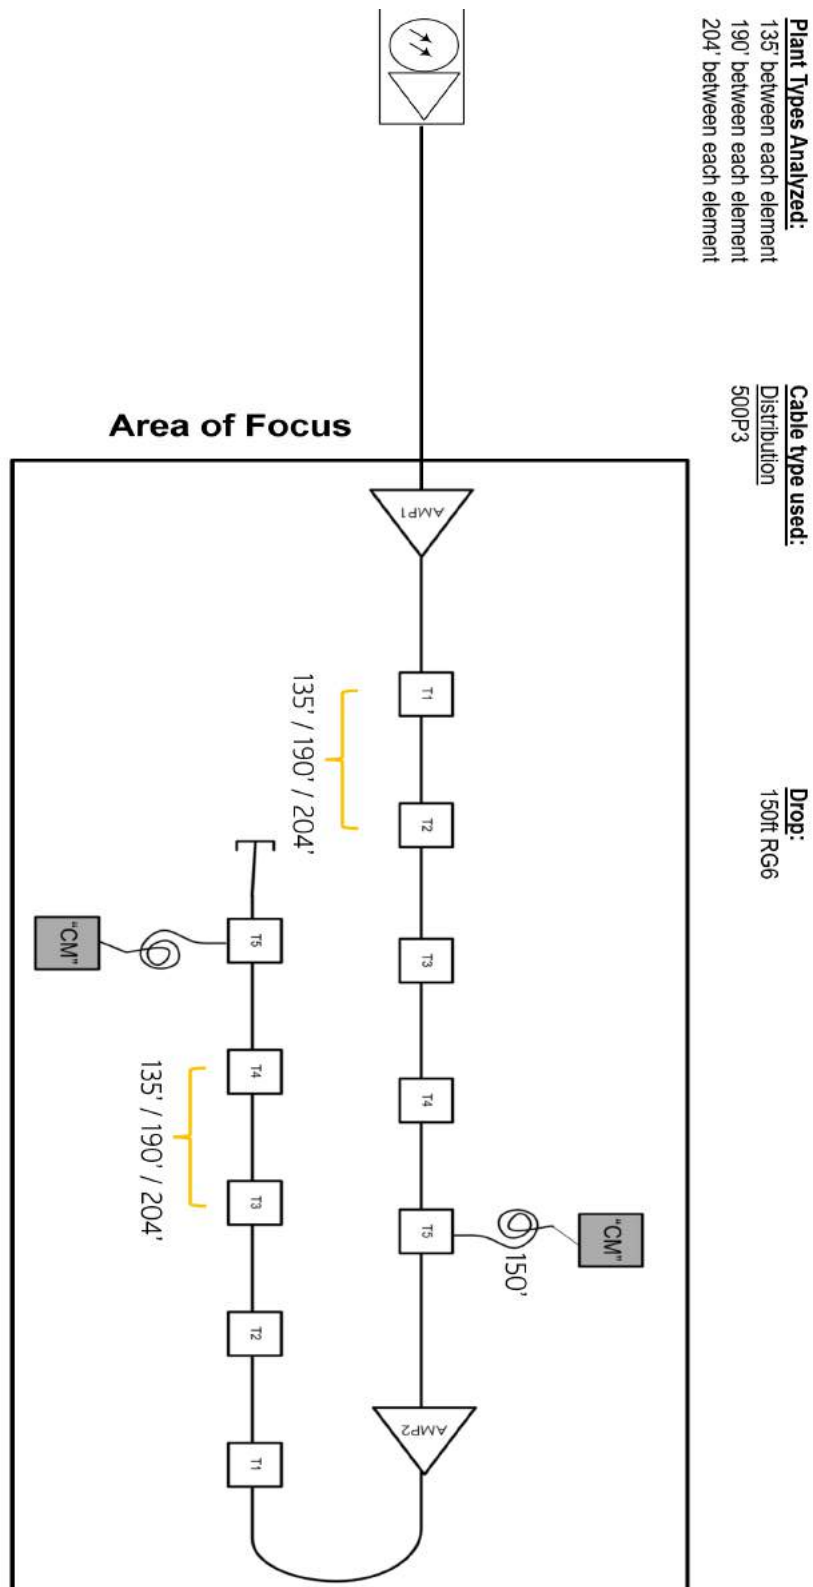
**Figure 4 – Plant Models**

Taking the span loss parameters above into consideration, the following plant models have been created for the analysis.

Note that the area of focus for this paper is in the last two spans of the N+2 plant, where the amplifiers are installed. The first span, between the node and the first amplifier has not been analyzed for performance. Instead the focus is on the input to the first amplifier since that will be the baseline for the system C/N.

Throughout this paper, each plant type will be referred to by its respective distance:

- 135' plant:** 135 feet between each plant element
- 190' plant:** 190 feet between each plant element
- 204' plant:** 204 feet between each plant element



**Figure 5 – Analyzed Plant Types**




#### 4.1. Traditional Node and Amplifier Outputs

Assuming 70 dBmV of TCP is available at the port of each active device, the following two options can be considered for extending the spectrum to 1.8 GHz:

1. Change the output tilt in a way to make it more ‘flat’
2. Introduce a step-down at a certain frequency, typically 1 GHz

Given the sensitivity of ‘legacy’ devices in the plant to RF level fluctuations, option 1 is typically avoided. Instead option 2 is typically considered by most operators in traditional plant design.

With that in mind, along with knowing that the majority of TCP is allocated on the higher portion of the spectrum (in this case 1.8 GHz), the following power spectral density (PSD) outputs have been assumed for traditional node and amplifier outputs:



**Figure 6 – Node and Amplifier Output PSD**

It can be seen from the figure above that distributed gain architecture (DGA) PSD does not have any step-downs throughout the spectrum, due to the addition of DGA amplifiers along the distribution path.

The raised levels from 258 MHz to 650 MHz should be mentioned in light of the comment above regarding legacy devices. Given that the low end of the spectrum has been raised by only 2.5 dB, the potential impact of this on legacy devices along the distribution path has been deemed insignificant.

## 4.2. Traditional Node and Amplifier Noise Figure (NF)

Depending on the type of amplifier deployed in the OSP and their respective internal splitting, the DS NF of traditional nodes and amplifiers can vary anywhere between 8 dB – 12 dB. In order to set a baseline for system C/N calculations, the following NF has been assumed:

**Table 1 – Node and Amplifier NF**

|                                                     |       |
|-----------------------------------------------------|-------|
| <b>Node &amp; Amplifier NF @ Device Port for DS</b> | 11 dB |
|-----------------------------------------------------|-------|

## 4.3. Modem (MDM) Transmit Power:

The following table has been referenced in the DOCSIS 4.0 specification for modem transmit power (Tx)


**Table 2 – MDM Tx Power/1.6 MHz**

| <b>Upstream Centre Frequency</b>            | <b>108 MHz</b> | <b>684 MHz</b> | <b>Spectral tilt (dB)</b> |
|---------------------------------------------|----------------|----------------|---------------------------|
| <b>Upstream Reference PSD (dBmV/1.6MHz)</b> | 33             | 43             | 10                        |

Converting the numbers above from 1.6 MHz reference PSD to 6.4 MHz equivalent numbers, the modem Tx power can be graphed as:

**Table 3 – MDM Tx Power/6.4 MHz**

| <b>Upstream Centre Frequency</b>            | <b>108 MHz</b> | <b>684 MHz</b> | <b>Spectral tilt (dB)</b> |
|---------------------------------------------|----------------|----------------|---------------------------|
| <b>Upstream Reference PSD (dBmV/6.4MHz)</b> | 39             | 49             | 10                        |



**Figure 7 – MDM Output Power PSD**

#### 4.4. Modulation Order vs. Power and C/N

In order to have a baseline for achievable modulation orders throughout the distribution plant, the below table from the DOCSIS 4.0 PHY specification has been utilized.

Note: although DOCSIS 4.0 modems are able to receive and demodulate signals as low as 16QAM with 16 dB of C/N and -30 dBmV/6 MHz, no values below 256QAM has been considered in this paper since modulation orders lower than 256QAM are typically deemed unacceptable by operators.

**Table 4 – MDM Tx Power/6.4 MHz**

| Constellation | C/N<br>(dB) | Rx Power/6 MHz<br>(dBmV) |
|---------------|-------------|--------------------------|
| 4kQAM         | 44          | -6                       |
| 2kQAM         | 40          | -9                       |
| 1kQAM         | 36          | -12                      |
| 512QAM        | 33          | -15                      |
| 256QAM        | 30          | -18                      |

## 5. DGA Amplifier Considerations

In order to implement booster or DGA amplifiers, a baseline for upstream (US) and downstream (DS) gains needs to be set. Given that these amplifiers are single stage with a fixed output, the output performance of the device is highly dependent on the input levels for upstream and downstream. This is defined by noise power ratio (NPR) for the US and carrier to interference noise ratio (CINR), as a function of input TCP.


For the designs carried out in this paper, due to lack of availability on the parameters mentioned above, instead only the US and DS amplifier's noise figure (NF) has been considered in the overall system performance considerations. The NF of booster/DGA amplifiers have been shown in the table below:

**Table 5 – DGA Amplifier DS and US NF**

|                                                                 |              |
|-----------------------------------------------------------------|--------------|
| <b>DGA/Booster Amplifier NF @ Device Port<br/>for US and DS</b> | <b>15 dB</b> |
|-----------------------------------------------------------------|--------------|

### 5.1. DS Gain


The following figure has been assumed for the DGA amplifier DS gain:



**Figure 8 – DGA Amplifier DS Gain**

## 5.2. US Gain

The following figure has been assumed for the DGA amplifier US gain:




**Figure 9 – DGA Amplifier US Gain**

## 6. Traditional Plant DS Results

The performance results in this section are used as a baseline for comparison. Before discussing the modem receive levels in each plant type, the amplifier contributions to system C/N for each scenario is calculated below.

### 135' Plant C/N:

Applying the 'traditional PSD' node and amplifier output in section 4.1 to the 135' plant model will result in the following Rx Power at the port of amplifiers:



**Figure 10 – 135' Amplifier Rx Power @ Amp. Port**


From the figure above the C/N contribution of the amplifiers at 1 GHz and 1.8 GHz can be calculated:

**Table 6 – 135' Traditional Amplifier Contributions to System CNR**

|                   |         |
|-------------------|---------|
| N+2 CNR @ 1 GHZ   | 58.7 dB |
| N+2 CNR @ 1.8 GHz | 59.4 dB |

### **190' Plant C/N:**

Applying the 'traditional PSD' node and amplifier output in section 4.1 to the 190' plant model will result in the following Rx Power at the port of amplifiers:



**Figure 11 – 190' Amplifier Rx Power @ Amp. Port**


From the figure above, the C/N contribution of the amplifiers at 1GHz and 1.8GHz can be calculated:

**Table 7 – 190' Traditional Amplifier Contributions to Ststem CNR**

|                   |         |
|-------------------|---------|
| N+2 CNR @ 1 GHZ   | 51.7 dB |
| N+2 CNR @ 1.8 GHz | 48.4 dB |

### **204 Plant C/N:**

Applying the ‘traditional PSD’ node and amplifier output in section 4.1 to the 204’ plant model will result in the following Rx Power at the port of amplifiers:



**Figure 12 – 204’ Amplifier Rx Power @ Amp. Port**

From the figure above the C/N contribution of the amplifiers at 1GHz and 1.8GHz can be calculated:

**Table 8 – 204’ Traditional Amplifier Contributions to System CNR**

|                   |         |
|-------------------|---------|
| N+2 CNR @ 1 GHZ   | 50 dB   |
| N+2 CNR @ 1.8 GHz | 45.6 dB |

**Observation:** From the results in Table 7 and 8, we can observe that although the input Rx power into the traditional amplifiers are below 11 dBmV/6 MHz, all the C/N’s are above the minimum required for 4kQAM. This is an optimistic assumption for 190’ and 204’ plant as the distribution network’s contribution the system C/N is very close to the numbers in Table 4. As noted in section 3.1, the starting C/N from the node plays a big role in the overall system C/N.

Knowing this and applying the ‘traditional PSD’ node and amplifier output discussed in section 4.1 to the plant models in section 4.0, will result in the modem receive levels (MDM Rx) below.

**Note:** as discussed in section 4, all the drops lengths throughout this paper are 150 feet of RG6 which can be considered a worst-case scenario.



## 6.1. 135' Plant MDM Rx Powers

The Rx power levels /6 MHz for each modem along the distribution line for the 135' plant has been demonstrated below:




Figure 13 – 135' Span 1 MDM Rx Power/6 MHz






Figure 14 – 135' Span 2 MDM Rx Power/6 MHz

## 6.2. 190' Plant MDM Rx Powers

The Rx power levels /6 MHz for each modem along the distribution line for the 190' plant has been demonstrated below:



**Figure 15 – 190' Span 1 MDM Rx Power/6 MHz**



**Figure 16 – 190' Span 2 MDM Rx Power/6 MHz**

### 6.3. 204' Plant MDM Rx Powers

The Rx power levels /6 MHz for each modem along the distribution line for the 204' plant has been demonstrated below:





Figure 17 – 204' Span 1 MDM Rx Power/6 MHz



**Figure 18 – 204' Span 2 MDM Rx Power/6 MHz**

**Observation:** in the 135' plant, all the modems are capable of receiving 4kQAM with traditional amplifiers only. On the contrary, it can be seen than 190' and 204' plant models struggle with achieving 4kQAM throughout the distribution network. It should also be emphasized that all of the analysis above was done using 150 feet of RG6 as the drops throughout the distribution plant. Reducing this length will improve the Rx levels at each MDM, improving the MER and achievable modulation order.

## 7. Booster Amplification DS Results

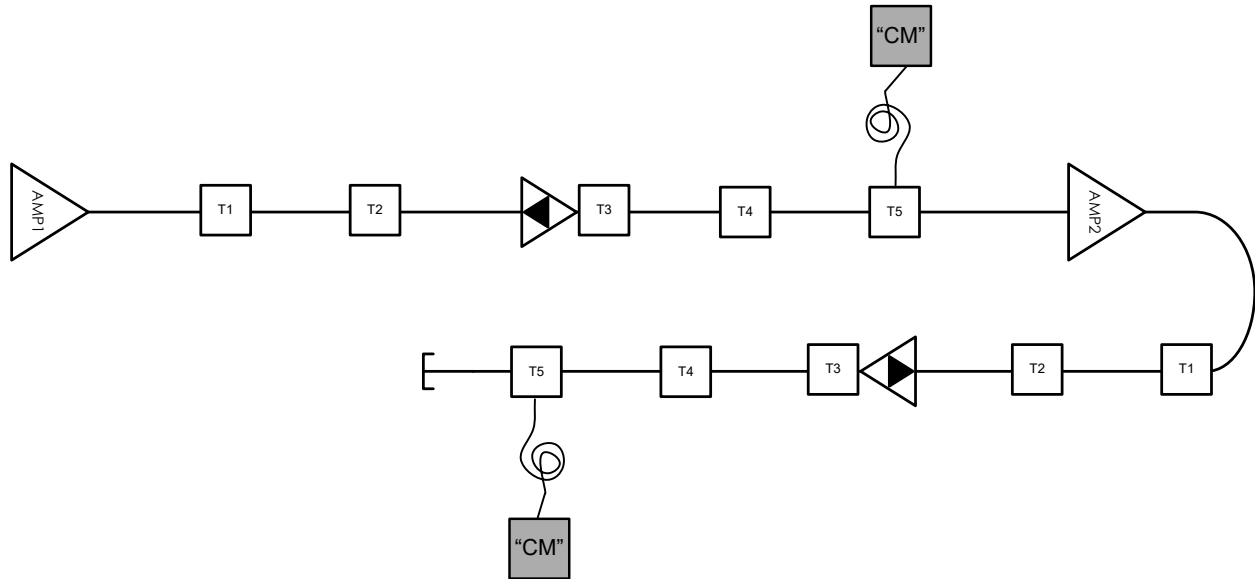
Based on the results demonstrated in the previous section, the 190' and 204' plant models struggle with achieving 4kQAM throughout the distribution network. These plant models could be prime candidates for adding booster amplifiers mid-span to not only increase the Rx power at each traditional amplifier input, but also to boost the MDM Rx levels in each span. Given that 135' plant is capable of achieving 4kQAM (Figures 13 and 14), no booster amplification has been considered for this plant model.

As discussed in section 4.1 and given the new additional mid-span gain from the DGA/booster amplifier, the need for a step downs in the output PSD is eliminated. This new output PSD results in a 2.5 dB lower TCP (69.8 dBmV vs. 67.3 dBmV), which can subsequently improve the output C/N of the node. As previously covered in section 3 this has not been considered in overall system C/N calculations.

**Note:** Based on the early prototype form factors of DGA/booster amplifiers, the final version of the product should approximately be the equal to the size of a mainline splitter. This can have immense benefits when it comes to ease of installation and access. In other words, as long as pedestals (PEDs) are installed for tap locations, an assumption has been made that DGA/booster amplifiers will fit in existing PEDs.

### 7.1. 190' and 204' Plant Design with Booster Amplifiers

The following design has been considered for the 190' and 204' plant types, with booster amplification mid-span:




**Figure 19 – 190' and 204' Plant with Mid-Span Booster Amplification**

In this new design with the added gain mid-span, the tap values after the booster amplifier must be increased to ensure reasonable Rx power at the modem. This will further improve the end-of-line Rx power because higher value taps have lower insertion loss values throughout the spectrum.

Additionally, as discussed in section 4, given the added 20 dB of gain at 1.8 GHz, there is no need for any step down at 1 GHz. DGA gain from Figure 8 has been applied to the node and amplifier outputs.


As a result, the new traditional amplifier Rx levels have been demonstrated below:

### **190' Plant Traditional Amplifier Rx Power with Booster Amplification:**



**Figure 20 – 190' Amplifier Rx Power @ Port with Booster Amplification**

### **204' Plant Traditional Amplifier Rx Power with Booster Amplification:**




**Figure 21 – 204' Amplifier Rx Power @ Port with Booster Amplification**

A point of concern would be adding two additional amplifiers, essentially taking the current N+2 design to N+4. The overall contributions of the amplifiers to the system C/N can be calculated for each case, demonstrated in the table below:

**Table 9 – Amplifier Contributions to System C/N**


|                                                | 190' Plant | 204' Plant |
|------------------------------------------------|------------|------------|
| <b>Amplifier Contributions to C/N @1 GHz</b>   | 61.8 dB    | 58.9 dB    |
| <b>Amplifier Contributions to C/N @1.8 GHz</b> | 58.7 dB    | 53 dB      |

An assumption can be made that the limiting factor in achieving each modulation order is the Rx power at the modem. Keeping that in mind, the figures below demonstrate MDM Rx power at each tap, with the addition of booster amplification mid-span:



**Figure 22 – 190' Span 1&2 MDM Rx Power/6MHz**

It can be observed that all the taps along the distribution line are now well above the 4kQAM threshold.



**Figure 23 – 204' Span 1&2 MDM Rx Power/6MHz**

Aside from ~200 MHz of Tap 5, it can be observed that all the taps along the distribution line are now above the 4kQAM threshold.

## 7.2. Booster Amplification Observations

Mid-span booster amplification seems to provide a viable option to increase the system C/N and subsequently, the overall system achievable modulation order. This is also assuming that the booster and DGA amplifiers have very low distortion characteristics, where they can be considered negligible.

## 8. DGA Design and DS Results

DGA design can be described as distributing the gain of traditional amplifiers along the path. This can potentially have the following benefits:

- Enhancing end-of-line performance
- Fully moving away from distortions and intermodulations and, as a result, achieving a noise-limited system
- Eliminating the need for having any step downs in node or amplifier output
- Reducing the output power of the node, resulting in improvement of the output MER
- Reduced power draw (covered in section 10)

Additionally, as discussed in section 7, the form factor of the DGA amplifiers are roughly the size of a mainline splitter. Assuming there are PED locations available for taps, DGA amplifiers should fit before or after the taps, depending on the design.

From an OSP design perspective, DGA can seem strange in comparison to designing a traditional plant. As demonstrated in sections 8.1, 8.2 and 8.3, the 'mid-span' tap values can vary anywhere from 17 to 20 dB taps, depending on where the DGA amplifier has been installed. This can present design challenges



for operators as new design methodologies have to be crafted in order to optimize plant performance. This will be discussed further in section 12.

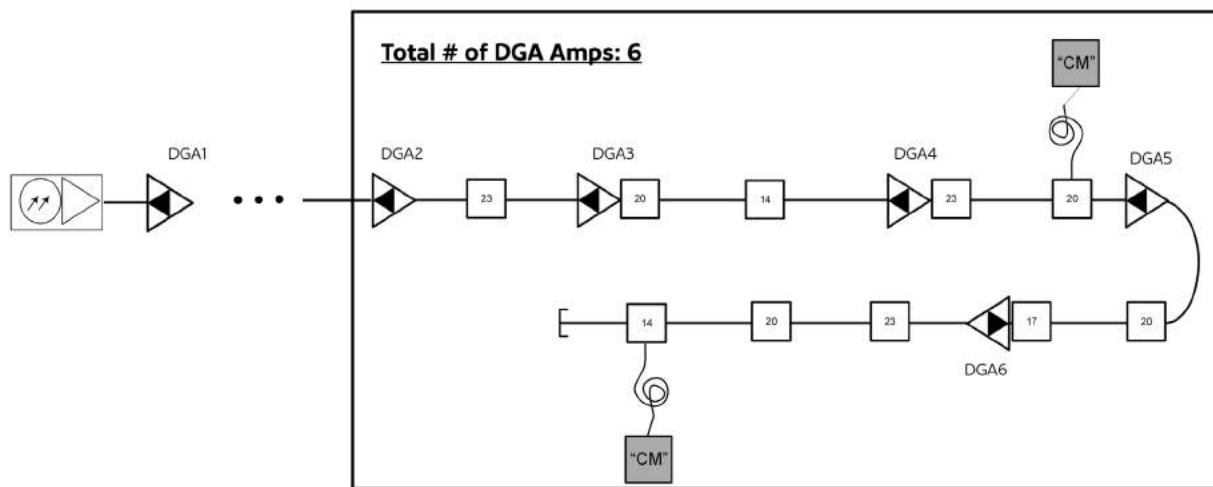
DGA designs might also seem counter intuitive to the concept of cascade reduction. As demonstrated in the following sections, the cascade length of the studied N+2 plants can increase up to 8. This cascade length can vary by approximately 2 amplifiers for each case analyzed, depending on the end-of-line performance expectations by the operator. For the purpose of this study, all plant types have been designed to achieve 4kQAM.

An important note to keep in mind is that the designs shown here are moving away from a unity gain design since no pads or equalizers were considered in this analysis. With traditional amplifiers, this can cause concern as distortions can accumulate quite rapidly when a system is not designed with unity gain in mind. Theoretically speaking, given the extremely low distortion characteristics of DGA amplifiers, it has been assumed that distortions will not result in degradation of signal quality at the end-of-line. This needs to be verified in the future as more products are available in this realm.

In the below sections, the performance of each plant type when converted to DGA, has been discussed.

### 8.1. 135' Plant Design and DS Results


The following design has been considered for 135' plant:



**Figure 24 – 135' DGA Conversion**

**Note:** in order to have a baseline, spans 1 and 2 have been kept the same in comparison to the ‘traditional’ plant spans.

As it can be seen, a previously N+2 plant has been converted to N+6. Although this might seem concerning at first, the Rx power at each DGA amplifier along with the amplifier contributions to the system C/N has been shown below:




**Figure 25 – 135’ Plant DGA Amplifiers’ Rx Power @ Ports**

From the figure above, calculating the overall amplifier contributions to system C/N will result in the following:


**Table 10 – 135’ DGA Amplifier Contributions to System C/N**

|                     | 135’ Plant |
|---------------------|------------|
| System C/N @1 GHz   | 61.8 dB    |
| System C/N @1.8 GHz | 58.7 dB    |

Knowing the above, it can be assumed that the limiting factor for plant performance would be the Rx power at the modems. The figures below demonstrate Rx power at each tap’s modem location for each span:



**Figure 26 – DGA 135' Span 1 MDM Rx Power/6MHz**

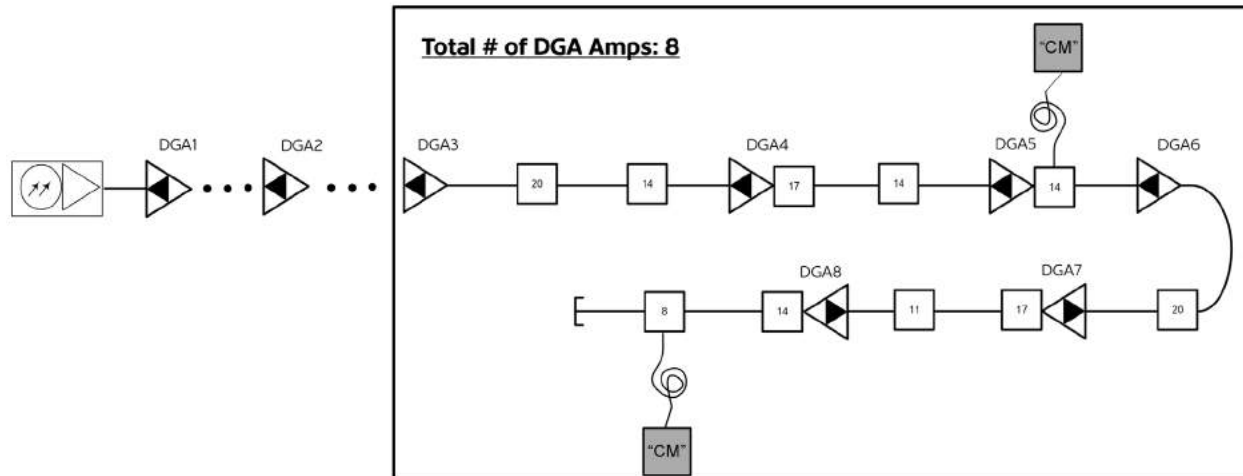


**Figure 27 – DGA 135' Span 2 MDM Rx Power/6MHz**

Although the 135' plant does not seem to benefit from DGA from an achievable modulation order perspective, it will benefit from the power reductions properties of DGA. This will be further discussed in section 10.


## 8.2. 190' Plant Design and DS Results

The following design has been considered for 190' plant:



**Figure 28 – 190' DGA Conversion**

The previous N+2 plant has been converted to N+8. The Rx power at each DGA amplifier has been shown below:



**Figure 29 – 190' Plant DGA Amplifiers' Rx Power @ Ports**

From the figure above, calculating the overall amplifier contributions to system C/N will result in the following:

**Table 11 – 190' DGA Amplifier Contributions to System C/N**

|                                                | 190' Plant |
|------------------------------------------------|------------|
| <b>Amplifier Contributions to CNR @1 GHz</b>   | 54 dB      |
| <b>Amplifier Contributions to CNR @1.8 GHz</b> | 54.5 dB    |

It can be seen that the overall system C/N can remain high, despite the fact that 8 amplifiers have been designed in cascade.

With that in mind, the modem Rx power still seems to be the limiting factor in the achievable modulation order. The figures below demonstrate Rx power at modems in each tap location:



**Figure 30 – DGA 190' Span 1 MDM Rx Power/6 MHz**

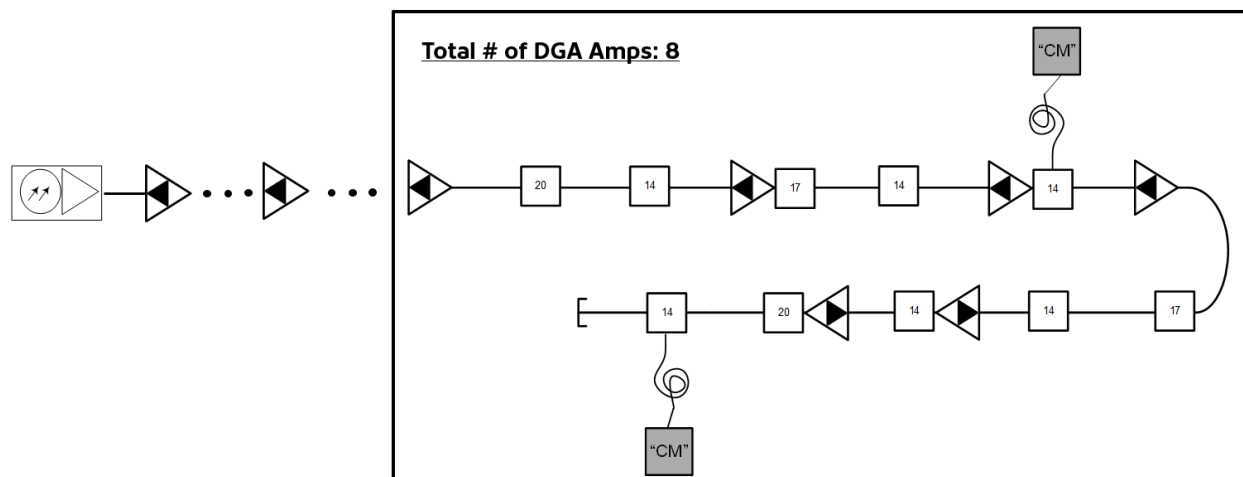


**Figure 31 – DGA 190' Span 2 MDM Rx Power/6 MHz**

It can be observed that all the Rx powers throughout the distribution plant are well above the minimum 4kQAM threshold.


### 8.3. 204' Plant Design and DS Results

The following design has been considered for 204' plant:



**Figure 32 – 204' DGA Conversion**

The previous N+2 plant has been converted to N+8. The Rx power at each DGA amplifier has been shown below:



**Figure 33 – 204’ Plant DGA Amplifiers’ Rx Power @ Ports**


From the figure above, calculating the overall amplifier contributions to system C/N will result in the following:

**Table 12 – 204’ DGA Amplifier Contributions to System C/N**

|                            | <b>204’ Plant</b> |
|----------------------------|-------------------|
| <b>System CNR @1 GHz</b>   | 54.2 dB           |
| <b>System CNR @1.8 GHz</b> | 53.5 dB           |

It can be seen that the overall system C/N remains high, despite the fact that 8 amplifiers have been designed in cascade.

With this in mind, the modem Rx power still seems to be the limiting factor in the achievable modulation order. The figures below demonstrate Rx power at modems in each tap location:



**Figure 34 – DGA 204' Span 1 MDM Rx Power/6 MHz**



**Figure 35 – DGA 204' Span 2 MDM Rx Power/6 MHz**



It can be observed that all the Rx powers throughout the distribution plant are above the minimum 4kQAM threshold.

## 8.4. DGA DS Design Observations


DGA appears to provide a high system C/N and end-of-line performance in each of the plant models analyzed. This is most visible in the 190' and 204' plant models analyzed. It can be observed that although the cascade length in the analyzed plant models were increased from +2 to +8, the achievable modulation orders were increased by roughly 2-3 orders of magnitude.

## 9. Upstream Analysis and Considerations

Given the complexity of US analysis in a system and due to noise funneling from amplifiers and modems, this paper has focused on the potential points of concern in a DGA plant, especially regarding 204' plant, given that it is the longest plant type analyzed.

It should be noted that no closed loop gain control has been considered in this analysis. It is assumed that all the modems in the distribution will be transmitting at maximum power in accordance to Figure 7, to determine any shortcomings in the upstream network.


Assuming that the modems sitting at each location will be transmitting with their maximum capability in the return path in accordance to Figure 32, the following Rx powers/6.4 MHz can be expected at the port of each DGA amplifier:



**Figure 36 – 204' Plant Return Path DGA Amplifiers' Rx Power @ Ports**

**Note:** In order to simplify the figure above, only data from tap values in parenthesis have been shown. These are modems that are subject to the highest amount of loss in the distribution plant. Furthermore, since the focus area of the analysis for this paper is in the distribution plant (DGA3-DGA8), no data from DGA1 and 2 have been shown in the figure above.

Although all the Rx power levels at the ports of each DGA amplifier seems sufficiently high, the primary point of concern is the return signal being subject high attenuations from long coaxial spans and higher insertion losses from low value taps. Assuming DGA amplifiers' return gain in Figure 9, the following Rx powers/6.4 MHz can be expected at the port of the node, in the 204' plant model:



**Figure 37 – 204' Plant Return Path Node Rx Power @ Port**

Given that the focus area of analysis for this paper is in the distribution portion of the plant (from DGA3 to DGA8), DGA1 and 2 have not been included in the figure above.


The figure above can raise concerns with regards to system performance in the US. This is especially the case because modems in DGA7 and 8 spans are the lowest common denominator, setting the limit for overall system performance in the return path.

The return path gain of DGA amplifiers is not the only point of concern. Given the DS DGA design in 204' plant, many 14 and 17 taps were used to ensure high DS Rx power levels at the MDM. Although this can improve the DS Rx power levels at MDM locations and subsequently increase the system's achievable modulation order in the DS, it will make the US performance suffer. For comparison, based on the tap data available today, typically 23 taps have 1.2-1.4 dB of loss in the legacy band. The same band will have 2 dB higher insertion loss in lower value taps such as 14 and 17 dB taps.

To show the significance of high insertion loss values from low value taps on the return path network, it can be assumed that higher value taps were used throughout the distribution network.

**Note:** In practice, this can be considered unrealistic in OSP designs, since increasing the tap values results in lower DS Rx power at the MDMs. This assumption is made simply to quantify the impact that low value taps can have in the overall system performance.

Assuming a lower insertion loss of 1.5 dB for the distribution taps, DGA8's new Rx Power at the node has been shown in comparison to Figure 37:




**Figure 38 – Last Span DGA-to-Node Level Comparison – 2dB Insertion Loss Taps vs. Regular Taps**

Although 1.5 dB of insertion loss may seem insignificant in traditional plant design, it makes a drastic difference in DGA. This is because in traditional HFC design, one or two low value taps may be installed in each distribution span. As shown in section 8.2 and 8.3, the number of low value taps can vary from 5-8 when converting a traditional N+2 environment.

In order to overcome the challenges in the US, two proposals are made in the following sections.

### 9.1. Higher MDM Transmit Levels in Lower US Splits


Knowing that the DOCSIS 4.0 MDM transmit channel set (TCS) will have 64.5 dBmV of TCP available and assuming that operators may not go to 684 MHz in the US, given the DS upper limit of 1.8 MHz, the ‘unused’ TCP from the upper frequencies in the return band can be re-allocated to the lower bands. This can result in an ‘up-lift’ of the transmit PSD of the MDM, which is demonstrated in the figure below:



**Figure 39 – Raised MDM Output Power PSD in Various Splits**

Increasing the transmit PSD may raise concerns regarding spurious emissions and fidelity requirements in accordance to the DOCSIS 4.0 specifications. This needs to be studied more extensively to ensure adherence to said specifications.

With that in mind, let us assume a 396 MHz split. Applying the 6 dB additional available power to DGA8 in Figure 37 will result in the following US Rx power at the port of the node:




**Figure 40 – Last Span DGA-to-Node Level Comparison – Raised MDM PSD**

It is visible that re-allocating the power from the unused 396 MHz - 684 MHz to the active 5 MHz to 396 MHz can approximately result in a 6 dB increase in US Rx power at the node without increasing the modem's TCP.


## 9.2. Higher Return Gain in DGA Amplifiers:

Let us assume the gain of the DGA return amplifier is increased by 2 dB at 108 MHz and by 1 dB at 684 MHz, resulting in the following figure:



**Figure 41 – High vs. Low Gain DGA Return Amplifier**

DGA8's new US Rx power at the node with the newly assumed gain has been demonstrated in the figure below:



**Figure 42 – Last Span DGA-to-Node Level Comparison – High Gain vs. Low Gain Return DGA Amplifier**

## 10. Power Draw

One of the most attractive concepts of DGA is reduced power draw in the OSP. Traditional plant design can be quite power hungry with 150 watt nodes, and amplifiers that draw anywhere from 30 to 60 watts.

Power supplies themselves can be quite challenging to deploy, given the reduced amount of available real estate for installing them. Due to this, the current power supplies installed in the OSP are expected to support the future technologies deployed by operators, which may seem very challenging

DGA can reduce the power consumption in the OSP drastically in comparison to traditional designs. In order to quantify this in the plant models analyzed in this paper, the following has been assumed:

- N+2 plant
- 4 outputs from each node
- 150 watt node
- 40 watt line extenders
- 7 watt DGA amplifier

The following table demonstrates the potential power saving in an N+2 plant.:

**Table 13 – Power Draw Comparisons**

| Traditional | Booster   | DGA 135' | DGA 204&190' |
|-------------|-----------|----------|--------------|
| 460 Watts   | 515 Watts | 320Watts | 375Watts     |
|             | ↑ 10%     | ↓ 30%    | ↓ 20%        |

Although the power draw in the booster amplification case has increased by 10%, it was demonstrated in section 7 that the performance at the end-of-line can increase by 2-3 orders of modulation. This is also assuming that future amplifiers will not be able to adjust power consumption based on their output power and utilized TCP.

It can be seen that DGA saves 20-30%, depending on the plant model. This is dependent on the number of DGA amplifiers used in each span to overcome the existing span losses.

This reduction in power draw from existing power supplies presents countless opportunities for operators to deploy other technologies in the access network, such as DAA, small cell and 5G.

## 11. Plant Reliability

Here we compare the availability of a traditional cable plant with the new distributed gain amplifier (DGA) system.

We assume that the failure modes of traditional amplifiers and DGAs are comparable, and one does not impact any more customers or impact any customers differently than the other. For example, customers on a branch are not impacted by the failure of an amplifier on another branch, and generally only the customers downstream of a failed amplifier are impacted by the failure of an amplifier.

Therefore, we can model these systems as simple series systems of replaceable components.

Further, we will define the components of each system as amplifiers and non-amplifiers.

The variables  $n_{\text{tamp}}$  and  $n_{\text{dga}}$  are respectively the number of traditional amplifiers and DGAs in the comparable systems

Given the systems are equivalent except for the number and type of amplifiers, we can define the availability of each of these systems as follows, for an arbitrary customer of the systems.

Traditional system availability:  $A_{\text{tsys}} = A_{\text{line}} * A_{\text{tamp}}^{n_{\text{tamp}}}$

DGA system availability:  $A_{\text{dgasys}} = A_{\text{line}} * A_{\text{dga}}^{n_{\text{dga}}}$

$A_{\text{line}}$  is the availability of the line system, everything but the amplifiers, which is the same in both architectures.  $A_{\text{tsys}}$  is the traditional system amplifier availability component, a series of amplifiers each with availability  $A_{\text{tamp}}$ .  $A_{\text{dgasys}}$  is the availability of the series of DGAs in the system, each with availability  $A_{\text{dga}}$ .

Now we compare the two systems. We want the new system to be at least as good as the old, so we have the constraint

$$A_{dgasys} \geq A_{tsys}$$

Or

$$A_{line} * A_{dga}^{n_{dga}} \geq A_{line} * A_{tamp}^{n_{tamp}}$$

So

$$A_{dga}^{n_{dga}} \geq A_{tamp}^{n_{tamp}}$$

Given that the number of amps in either system is an integer, and  $n_{dga} = n_{tamp} + n$  for some integer  $n$  greater than 0, we can rewrite the previous equation as

$$A_{dga}^{n_{tamp}+n} \geq A_{tamp}^{n_{tamp}}$$

And then taking the  $n_{tamp} + n$  root of both sides, with acknowledgement that the variables are bounded positive, and the availabilities are between 0 and 1, we get

$$A_{dga} \geq A_{tamp}^{n_{tamp} / (n_{tamp} + n)}$$

We refer later to this above equation as the amplifier relation. Considering the architectures analyzed:

$$\text{If } n_{tamp} = 2, n = 6, n_{tamp} / (n_{tamp} + n) = 0.25$$

$$\text{If } n_{tamp} = 2, n = 5, n_{tamp} / (n_{tamp} + n) = 0.29$$

$$\text{If } n_{tamp} = 2, n = 4, n_{tamp} / (n_{tamp} + n) = 0.33$$

The tighter constraint is 0.25, and so a goal to reach is

$$A_{dga} \geq A_{tamp}^{0.25}$$

But most plant designs will only use 6 DGAs so the more common constraint will be


$$A_{dga} \geq A_{tamp}^{0.33}$$

If the probability that the above two equations are each true is greater than 50%, then odds are the new architecture will perform better on average in a large sample.


Now let's look at some comparisons based on amp availability values.

We can assume reasonably that  $A_{tamp}$  ranges from 0.999 to 0.99999 given service performance; with some data collection we can narrow it down further, and even find estimates based on use conditions, environment, etc. But based on this broad range of estimates, setting the inequality in the amplifier relation equation to an equality to see the worst case for the DGAs, we find that  $A_{dga}$  relates to  $A_{tamp}$  as in the figures below. We show results as dots for two deployments: 6 DGAs and 8 DGAs. The line of equality is added as a solid line for reference. The first graph shows availability in linear scale, while the second shows unavailability (1-availability) in log scale ( $U_{dga}$  and  $U_{tamp}$  respectively).





**Figure 43 – DGA Availability**



**Figure 44 – DGA Un-availability**

**Table 4 – Unavailability Comparison for DGAs for Two Example Architectures versus Traditional Amplifiers**

| $U_{dga}$ , 8 DGAs | $U_{dga}$ , 6 DGAs | $U_{tamp}$ | %diff, 8 DGAs | %diff, 6 DGAs |
|--------------------|--------------------|------------|---------------|---------------|
| 0.0002501          | 0.0003334          | 0.0010000  | 74.991%       | 66.656%       |
| 0.0002251          | 0.0003001          | 0.0009000  | 74.992%       | 66.657%       |
| 0.0002001          | 0.0002667          | 0.0008000  | 74.992%       | 66.658%       |
| 0.0001750          | 0.0002334          | 0.0007000  | 74.993%       | 66.659%       |
| 0.0001500          | 0.0002000          | 0.0006000  | 74.994%       | 66.660%       |

| U_dga, 8 DGAs | U_dga, 6 DGAs | U_tamp    | %diff, 8 DGAs | %diff, 6 DGAs |
|---------------|---------------|-----------|---------------|---------------|
| 0.0001250     | 0.0001667     | 0.0005000 | 74.995%       | 66.661%       |
| 0.0001000     | 0.0001333     | 0.0004000 | 74.996%       | 66.663%       |
| 0.0000750     | 0.0001000     | 0.0003000 | 74.997%       | 66.664%       |
| 0.0000500     | 0.0000667     | 0.0002000 | 74.998%       | 66.665%       |
| 0.0000250     | 0.0000333     | 0.0001000 | 74.999%       | 66.666%       |
| 0.0000225     | 0.0000300     | 0.0000900 | 74.999%       | 66.666%       |
| 0.0000200     | 0.0000267     | 0.0000800 | 74.999%       | 66.666%       |
| 0.0000175     | 0.0000233     | 0.0000700 | 74.999%       | 66.666%       |
| 0.0000150     | 0.0000200     | 0.0000600 | 74.999%       | 66.666%       |
| 0.0000125     | 0.0000167     | 0.0000500 | 75.000%       | 66.666%       |
| 0.0000100     | 0.0000133     | 0.0000400 | 75.000%       | 66.667%       |
| 0.0000075     | 0.0000100     | 0.0000300 | 75.000%       | 66.667%       |
| 0.0000050     | 0.0000067     | 0.0000200 | 75.000%       | 66.667%       |
| 0.0000025     | 0.0000033     | 0.0000100 | 75.000%       | 66.667%       |

From Figure 43, it appears that the DGAs need to have a much higher availability than the traditional amplifiers. But looking at Figure 44, from an unavailability (downtime) perspective, the difference is not unlikely to be achieved. Table 4 shows the values used to plot Figure 44. The percent difference in unavailability (%diff) is calculated as  $(U_{tamp} - U_{dga})/U_{tamp}$ .

For comparative perspective, the DGA needs to have about 75% less unavailability than the traditional amplifier over its useful life if 8 are used. But if just 6 are used, then just 67% less unavailability is needed. Think of this percentage of unavailability as a reduction in downtime overall. While a  $\frac{3}{4}$  reduction in downtime might seem aggressive, these DGAs are much simpler, with newer components, so have a chance to beat that mark if designed for reliability. Recall that in most cases only 6 DGAs are needed, so most cases need just a  $\frac{2}{3}$  reduction.

An additional consideration too is that, if DGAs have significantly higher availability, it is not likely because they are significantly more repairable, but rather because their rate of occurrence of failures over the same lifetime of the original amplifiers is much lower. This means DGAs are likely to have much longer useful lifetimes, and therefore may further reduce lifetime costs for providing service. However, highly accelerated life testing should be conducted to verify this assertion, and to provide evidence of the seemingly aggressive availability and reliability targets.

Note that the same equations describe the relationships for reliability as well as availability, as these are series systems. But as this is a repairable system, availability is the measure that makes more sense for the system. Fortunately, the system availability is an important contributing factor to the service availability which is important for customer experience. Reliability of the service from a user experience may be important as well.

## 12. Future Considerations

This section summarizes and discusses the points of concern that were brought up in the previous sections of this paper.

## 12.1. Unity Gain, Distortions and Cascade Limits

One of the primary assumptions of this paper was that DGA amplifiers' distortions accumulate less quickly than the current amplifiers deployed by operators. This is primarily due to the simplicity and the low gain characteristics of these amplifiers, being single stage, with a fixed gain and tilt.


As mentioned in section 3.3, unity gain has been one of the primary design focuses for OSP in the past. Operators have adjusted their networks by adjusting node and amplifier outputs to balance both noise and distortions. Moving away from unity gain can raise concerns, especially regarding distortions. Knowing that almost all carriers deployed in the spectrum in the future will be digital, the distortion products can be summed into CIN, which can be translated to increase in noise level that should be considered in the system C/N. Due to lack of availability of data in this realm, this was not considered in this paper but it is something that needs to be studied extensively.

This subject needs to be studied further in the future when more products are available in this realm, to ensure that CIN products will not decrease the overall signal quality.

It should be noted that unity gain design can theoretically be achieved with DGA, assuming adjustments can be made in DGA amplifiers. This can also increase the number of DGA amplifiers needed in comparison to a non-unity gain design. It is worth noting that unity gain designs with DGA should be easier to achieve in green-field, in comparison to drop-in upgrades in brown-field applications. This can be achieved by simply balancing span losses and gain of DGA amplifiers, as currently done in traditional plant design.

The increased number of DGA amplifiers itself can raise concerns as well. Afterall, operators have been reducing cascades by pushing fibre deeper into the distribution plant. As demonstrated in section 8, when converting an N+2 plant, the number of DGA amplifiers can vary anywhere from 6-8, when designing for optimal performance (4kQAM).

When discussing the potential maximum number of DGA amplifiers in cascade, it requires a fine balance between the available spectrum, system C/N and distortions (CIN). The figure below, extracted from Broadband Cable Access Networks by David Lafarge and James Farmer, demonstrates this perfectly:



**Figure 45 – Relationship between Cascade, Noise and Distortion**

## 12.2. US Gain and Performance

As discussed in section 9, the US system performance can be a major point of concern, especially in longer span (higher span-loss) plant types. This is due to the increase amount of loss that the signal is subjected to. The higher loss is due to the longer coaxial distances along with higher insertion losses in low value taps.

The following proposals were made in sections 9.1 and 9.2:

- Raised output PSD's at the MDM
- Increased return gain in DGA amplifiers

The two proposals have to be studied individually and in combination to determine the feasibility of them being implemented by the vendor community.

## 12.3. Design Standards

DGA presents a major shift in how access networks are designed. Given that OSP designs have remained more or less the same in the previous decades, such a drastic shift in design could be a multi-year endeavor.

A slower and incremental implementation of DGA is something to be considered. As MSO's reduce cascade lengths in the OSP while pushing fibre deeper, booster amplifiers can be implemented in sections of the plant that struggle with the current spacing, especially when upgrading the available spectrum to 1.8 GHz. This can also help with not having to potentially deploy DGA in N+3 or N+4 architectures. This may result in 15+ amplifier cascades which can present difficulties regarding plant maintenance and performance in the future.

### 12.4. 3 GHz

DGA presents exciting insights into what access architecture and designs could look like when contemplating 3 GHz spectrum expansions. In 2019, a paper was published under the title “Blueprint for 3 GHz, 25 Gbps DOCSIS” by John T Chapman, Hang Jin, Thushara Hewavithana and Rainer Hillermeier, which covers this topic in great detail. As MSO’s reduce cascades and reach passive networks (N+0) in the future, DGA can be the answer to achieving 3GHz of available spectrum and 25 Gbps.

## 13. Conclusion

This paper discussed how booster amplification and DGA implementations can improve performance in the analyzed plant models. A comparison for overall system performance and C/N contributions from amplifiers were discussed between a traditional N+2 plant models, N+2 models with the introduction of booster amplifiers mid-span, and DGA converted versions of them.

From a downstream performance perspective, it was observed that the analyzed plant models in a DGA system can achieve much higher orders of modulation in comparison to the traditional plant model. This is despite increasing the cascade length from +2 to +6 or +8, depending on the spacing. This was especially visible in longer plant models that have been ‘stretched’ to their current spacing limits.

It was also demonstrated that DGA can substantially reduce the power draw in the OSP, leaving headroom for future technologies to be deployed in the access network. This is one of the most attractive points of a DGA system, given the current limitations and challenges that operators face regarding plant powering.

Furthermore, the form-factor of the DGA amplifiers alleviate a number of concerns regarding available real estate for installing traditional amplifiers and PEDs. High gain amplifiers require large PEDs to accommodate their large form-factors, along with cooling requirements. DGA amplifiers’ form-factors are more comparable to main-line splitters, making them extremely convenient for installation in the OSP. They can be installed in almost any existing PED in the access network.

It was also demonstrated that a DGA plant can be as reliable as a traditional one, despite the increase in the number of amplifiers in a distribution run.

As attractive as the items mentioned above may be, DGA can be seen as a major shift in how access networks are designed. Various challenges were also discussed throughout this paper that require more in-depth research. Upstream gain and performance of a DGA system, maximum cascade length and unity gain designs are some of the points of concern.

Although this concept is in its infancy, it does seem to offer extremely attractive solutions to some of the major challenges that operators may face when upgrading their networks to 1.8 GHz and beyond. It would be worth the effort to analyze the challenges identified, to ensure they can be alleviated.

# Abbreviations

|        |                                                 |
|--------|-------------------------------------------------|
| C/N    | carrier to noise ratio                          |
| CIN    | carrier to intermodulation noise                |
| CINR   | carrier to interface noise ratio                |
| CSO    | composite second order distortion               |
| CTB    | composite triple beat distortion                |
| dB     | decibels                                        |
| dBmV   | decibels relative to one millivolt              |
| DAA    | distributed access architecture                 |
| DGA    | distributed gain architecture                   |
| DOCSIS | data over cable service interface specification |
| DS     | downstream                                      |
| ESD    | extended spectrum DOCSIS                        |
| GHz    | gigahertz                                       |
| HFC    | hybrid fibre-coax                               |
| ISBE   | International Society of Broadband Experts      |
| MDM    | modem                                           |
| MER    | modulation error ratio                          |
| MHz    | megahertz                                       |
| NF     | noise figure                                    |
| NPR    | noise power ratio                               |
| OFDM   | orthogonal frequency division multiplexing      |
| OSP    | outside plant                                   |
| PED    | pedestal                                        |
| PoE    | point of entry                                  |
| PSD    | power spectral density                          |
| QAM    | quadrature amplitude modulation                 |
| RF     | radio frequency                                 |
| Rx     | receive                                         |
| SCTE   | Society of Cable Telecommunications Engineers   |
| TCS    | transmit channel set                            |
| TCP    | total composite power                           |
| Tx     | transmit                                        |
| US     | upstream                                        |

# Bibliography & References

Data-Over-Cable Service Interface Specification DOCSIS 4.0 – *Physical Layer Specification CM-SP-PHYv4.0*

Broadband Cable Access Networks – The HFC Plant, David Lafarge and James Farmer

# Closed Loop Capacity Optimization for Extended Spectrum DOCSIS

A Technical Paper prepared for SCTE•ISBE by

**Dr. Thushara Hewavithana**

Senior Architect, Network Platform Group  
Intel Corporation  
5000 West Chandler Blvd, Chandler, AZ 85226  
602-245-1468  
Thushara.hewavithana@intel.com

**Dr. Rainer Strobel**

Lilienthalstr. 16, 85579 Neubibrg  
rstrobel@maxlinear.com

**Nader Foroughi**

Senior Network Architect, Access Architecture & Technology  
Shaw Communications Inc  
2728 Hopewell Place NE. Calgary, AB T1Y7J7

# Table of Contents

| Title                                                                           | Page Number |
|---------------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                            | 4           |
| 2. Technological and Operational Challenges with Extended Sepctrum DOCSIS ..... | 5           |
| 2.1. Plant Spacing and Drop-In Upgrades .....                                   | 5           |
| 2.2. Total Composite Power (TCP) .....                                          | 5           |
| 2.3. Passive Network Updates – Taps and Other Connectors .....                  | 6           |
| 2.4. Plant Model .....                                                          | 6           |
| 2.4.1. N+0 vs N+X.....                                                          | 8           |
| 2.4.2. Noise .....                                                              | 8           |
| 2.4.3. Distortion .....                                                         | 9           |
| 2.4.4. Designing a Noise-Limited System.....                                    | 9           |
| 2.5. Typical Plant Topologies .....                                             | 10          |
| 2.6. Problem Definition .....                                                   | 11          |
| 3. Theoretical Framework for Closed Loop Throughput Optimization .....          | 11          |
| 3.1. Node + 0 Network .....                                                     | 11          |
| 3.2. Extention to Node + X, $X > 0$ , Networks .....                            | 14          |
| 3.3. Algorithm Description .....                                                | 16          |
| 4. System Level Solution for Closed Loop Optimization .....                     | 16          |
| 4.1. Full Spectrum allocated to each CM .....                                   | 17          |
| 4.2. Channel Stacking .....                                                     | 17          |
| 4.3. Implementation Considerations.....                                         | 17          |
| 4.4. Create Headroom in Power Budget for Soft Flexible MAC Architecture .....   | 18          |
| 5. Simulation Results and Discussion .....                                      | 19          |
| 5.1. Full Optimization for Node + 0 and Node + 4.....                           | 19          |
| 5.1.1. Node + 0 Network, Mid-Split .....                                        | 19          |
| 5.1.2. Node + 4 Network, Mid-Split .....                                        | 20          |
| 5.1.3. Node + 0 Network, High-Split .....                                       | 22          |
| 5.1.4. Node + 4 Network High-Split .....                                        | 23          |
| 5.2. With Staggered Channel Allocation – Channel Stacking.....                  | 24          |
| 5.2.1. Node + 0 Network Mid-Split .....                                         | 27          |
| 5.2.2. Node + 4 Network Mid-Split .....                                         | 27          |
| 5.2.3. Node + 0 Network High-Split .....                                        | 28          |
| 5.2.4. Node + 4 Network High-Split .....                                        | 28          |
| 6. Conclusion and Future Work.....                                              | 29          |
| Abbreviations .....                                                             | 29          |
| Bibliography & References.....                                                  | 30          |

## List of Figures

| Title                                                                                                 | Page Number |
|-------------------------------------------------------------------------------------------------------|-------------|
| Figure 1 – Agrregate rates of different access network topologies by standard introduction year ..... | 4           |
| Figure 2 – Amplifier Nonlinear Distortion .....                                                       | 6           |
| Figure 3 – Trunk Span Losses.....                                                                     | 7           |
| Figure 4 – Distribution Span Losses .....                                                             | 8           |
| Figure 5 – Signal Level Balanced Between Noise and Distortion.....                                    | 10          |
| Figure 6 – Node + 0 Passive HFC Plant Topology .....                                                  | 10          |
| Figure 7 – Node + 4 Cascade Plant Topology .....                                                      | 11          |



|                                                                                                                                        |    |
|----------------------------------------------------------------------------------------------------------------------------------------|----|
| Figure 8 – Distortion evaluation from amplifier circuit model .....                                                                    | 13 |
| Figure 9 – Node+X topology with Power Spectrum Optimization .....                                                                      | 15 |
| Figure 10 – Downstream Transmit Power allocation as seen from Node Interface C .....                                                   | 18 |
| Figure 11 – Rate for flat RX (blue), flat TX (red) and optimal (yellow) power allocation .....                                         | 20 |
| Figure 12 – Attenuation of the cable sections (left) and signal and noise PSDs at the last Tap (right).....                            | 20 |
| Figure 13 – Node + 4 mid split SNR at Tap 3 (left) and bit allocation at Tap 3 (right).....                                            | 21 |
| Figure 14 – Node + 4 mid-split Data rates for tilted TX- flat RX PSD (blue), flat TX PSD (red), and<br>optimized PSD (yellow) .....    | 21 |
| Figure 15 – Node + 4 mid-split Transmit PSDs and TCPs for different allocation schemes .....                                           | 22 |
| Figure 16 – Node + 0 high-split Data rates for flat RX (blue), flat TX (red), and optimized (yellow) and<br>corresponding TCP .....    | 22 |
| Figure 17 – Node + 0 high-split Transmit PSDs for different allocation schemes .....                                                   | 23 |
| Figure 18 – Node + 4 high-split Data rates for flat RX (blue), flat TX (red) and optimized (yellow) and<br>the corresponding TCP ..... | 23 |
| Figure 19 – Node + 4 high-split Transmit PSDs for different allocation schemes .....                                                   | 24 |
| Figure 20 – Channel to CM allocation for PSD optimization .....                                                                        | 24 |
| Figure 21 – Power allocation vs Data rates .....                                                                                       | 25 |
| Figure 22 – Channel stacking for two service groups.....                                                                               | 25 |
| Figure 23 – Default channel stacking scheme for two service groups.....                                                                | 26 |
| Figure 24 – Node + 0 mid-split Power allocation and Rate for channel stacking .....                                                    | 27 |
| Figure 25 – Node + 4 mid-split power allocation and Rate for channel stacking.....                                                     | 27 |
| Figure 26 – Node + 0 high-split power allocation and Rate for channel stacking.....                                                    | 28 |
| Figure 27 – Node + 4 high-split power allocation and Rate for channel stacking.....                                                    | 28 |


## List of Tables

| <b>Title</b>                                   | <b>Page Number</b> |
|------------------------------------------------|--------------------|
| Table 1 – Optimization algorithm summary ..... | 16                 |

## 1. Introduction

Achievable data rates in the fixed access network keep increasing. While passive optical networks (PON) move to 25 or 50 Gbit/s with IEEE 802.3ca [1] and MGfast [2] targets 10 Gbit/s aggregated point-to-point rate over twisted pair, it is time to evaluate HFC technology as a successor for 10 Gbit/s capable DOCSIS, using full duplex or 1.8 GHz bandwidth [4].

Figure 1 compares the data rate trends for different access technologies. DSL, as a point-to-point technology is at lower rates, but with a higher growth rate. While DOCSIS and PON, both shared medium technologies, follow a similar trend with lower growth rate of the aggregated rate, which is compensated by reducing the number of subscribers sharing the bandwidth as an additional measure.



**Figure 1 – Aggregate rates of different access network topologies by standard introduction year**

From Figure 1, aggregated data rates around 30 Gbit/s are a competitive choice for a future DOCSIS generation, which is herein called extended spectrum DOCSIS (ESD). This will allow 20-25 Gbit/s downstream (DS) and 5-10 Gbit/s upstream (US) rates, which is comparable to a single 25G PON wavelength service that is shown in [5] to serve future access network requirements. Following the arguments of [6], this will allow for 10 Gbit/s services and cover the bandwidth growth predicted by Nielsen's law [7].

The cable industry has recognized ESD as a viable path to extend competitiveness of DOCSIS network at a fraction of cost compared to fiber deployments going forward. Under the 10G DOCSIS initiative, 1.8 GHz frequency division duplexing (FDD) DOCSIS and 1.2 GHz full duplex DOCSIS options have been included in DOCSIS 4.0 as two possible ways of getting to 10 Gbit/s node throughput, enabling low single digit Gbit/s services. We can consider 1.8 GHz FDD to be an intermediate step to get to the 3 GHz ESD and 10 Gbit/s services. In this paper, we will focus more on 1.8 GHz ESD when describing algorithms and evaluating performance results. Nonetheless we will maintain the forward compatibility of our algorithms for a future 3 GHz ESD solution.

One of the key enablers of ESD is the advancement of power amplifier (PA) technology that can support multi GHz transmit signal. However, total composite power (TCP) of these PAs does not scale with the

increased spectrum beyond 1218 MHz. Therefore, the ESD communication system is limited in its transmit power. Optimal allocation of available transmit power and appropriate bit-loading (profile definition) is needed to get the maximum capacity out of the network.

In this paper, we outline a framework for closed loop optimization of the capacity of ESD systems subjected to the TCP constraint mentioned above. Cable modems provide the node with channel estimate and signal to noise ratio estimate data and the intelligent node uses this data to calculate the optimal power allocation and bit-loading for the downstream. This can be made part of the profile management application running on a virtual Cable Modem Termination System (vCMTS). We have shown that combining careful allocation of channels, closed loop optimization of transmit power, and adaptive bit-loading achieves considerable gains in data rate and reduction in TCP for network topologies currently present in MSO networks.

## **2. Technological and Operational Challenges with Extended Spectrum DOCSIS**

As MSO's expand their spectrums to 1.8 GHz and beyond, the expectation for the achievable modulation order at the top end of the spectrum is reduced. In traditional plant, being 750 MHz or even 1 GHz, most MSOs expect 4096 QAM to be achievable by each orthogonal frequency division multiplexing (OFDM) block deployed there. But the same is not true for 1.8 GHz and beyond. There are many reasons why that is the case but the primary one would be the current spacing that outside plant (OSP) is designed to. Due to this, any method to increase the achievable throughput in the network is highly sought after.

This section explores the technological and operational challenges of upgrading the plant to 1.8 GHz.


### **2.1. Plant Spacing and Drop-In Upgrades**

Most OSP architectures are designed to 750 MHz and 'stretched' to 1 GHz. As a result, most of the RF power in plant actives, including nodes and amplifiers, has been utilized to overcome the existing span losses. Span loss is defined as the total insertion loss of all the elements in an HFC span, measured in dB. This includes all the plant passives such as splitters and couplers. Although the span losses today are manageable with the current amplifier gains, they will certainly become a major point of concern when the spectrum is expanded to higher frequencies. As an example, a span loss of 35 dB in a traditional 1 GHz plant equates to 49 dB in 1.8 GHz.

Although plant re-spacing is always an option, it can be extremely costly and as a result, operators will rely on the expanded power of amplifier gain chips to overcome span losses.

### **2.2. Total Composite Power (TCP)**

It is generally understood that the 1.8 GHz amplifiers will have ~75 dBmV of total composite power (TCP) available, which is roughly the same power available in current 1.2 GHz devices. With that been said, not all of this power is available for use. As an example, Figure 2 demonstrates the trade-off between Modulation Error Ratio (MER) and TCP utilized.



**Figure 2 – Amplifier Nonlinear Distortion**

As a general rule of thumb, 3 dB of back-off is needed to achieve 40 dB+ MER, which is typically what operators aim for. Along with that, the internal loss of the active device has to be accounted for, which is typically 2 dB. To summarize, there is a total of 70 dBmV of power to be utilized at the port of each active device. This can be a concern for operators given that 65-68 dBmV of TCP has already been allocated to overcome the span losses in the ‘traditional’ plant.

### 2.3. Passive Network Updates – Taps and Other Connectors

Taps and passives can be another point of concern when upgrading the OSP to 1.8 GHz. Traditionally, most operators have relied on face-plate upgrades to expand the spectrum range of plant taps and passives. This is generally accepted as a faster and more cost-effective method to upgrade the available bandwidth of taps.

Unfortunately, this might not be the case with 1.8 GHz upgrades. Based on the research and the information released by the vendor community, a face-plate upgrade of the current 1 GHz taps can potentially expand the bandwidth to ~1.6 GHz. it should also be noted that this is a best effort.

This can be a concern given the uncertainty of the maximum available bandwidth in the plant. As a result, it is generally accepted that taps and passives have to be swapped out for 1.8 GHz version. Given that the entire housing of the tap has to be swapped out as a part of this effort, most of the taps being developed will have housings that can support up to 3 GHz with future face-plate upgrades, future proofing the plant for 3 GHz upgrades.

### 2.4. Plant Model


In order to encompass most of the hybrid-fiber-coax (HFC) architectures deployed, the following plant models and assumptions were considered for this analysis.

Note: Trunk spans are defined as spans that are untapped. Distribution spans are tapped. Both trunk and distribution span losses include all other passive elements' insertion losses, such as splitters and couplers.

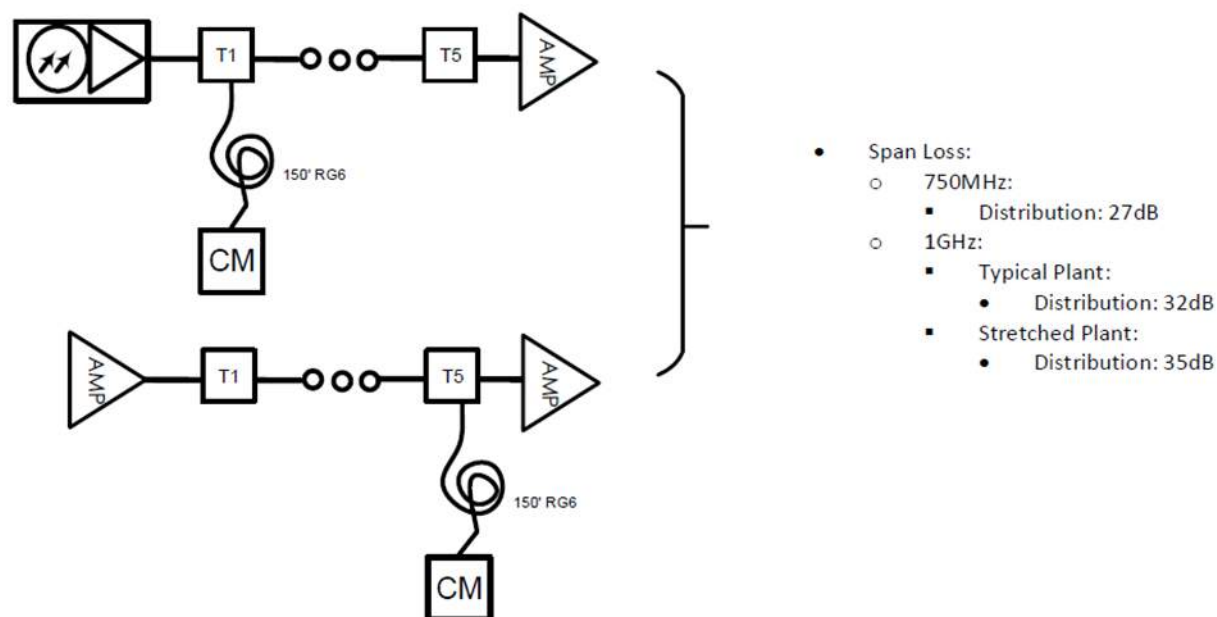
Assumptions:

- Modem:
  - Point of entry (PoE) device
- Drop:
  - Cable: RG6
  - Length: 150 feet
- Number of taps in each span:
  - 5
- Distribution span losses at 1 GHz:
  - Typical plant: 35 dB
  - Stretched plant: 37 dB
- Trunk span losses at 1 GHz:
  - Typical plant: 32 dB
  - Stretched plant: 35 dB

Figure 3 and Figure 4 summarize the parameters above.



**Figure 3 – Trunk Span Losses**



**Figure 4 – Distribution Span Losses**

### **2.4.1. N+0 vs N+X**

HFC plant can be divided into two categories: passive and cascaded.

- Passive plant (N+0): where no amplifiers are used after the node
- Cascaded plant (N+X): where amplifiers are used to boost the signal multiple times to the end-of-line

When designing an N+0 plant, the main point of concern is the output power of the node. Assuming that we are operating in a distributed access architecture (DAA) plant, the primary drivers for the plant quality would be the modulation error ratio (MER) of the DAA device. Since no amplifiers are used to boost the signal, no noise or distortion is added to the primary signal being generated by the DAA device.

On the contrary, when designing a cascaded plant, the following can be a concern:

- Amplifier noise contribution
- Amplifier distortion contribution

### **2.4.2. Noise**

Designing a cascaded system for optimal carrier to noise is always a big priority for an operator. One of the biggest contributors in system design is the receive power (Rx Power) at the amplifier, given that it is one of the primary drivers for the overall system carrier to noise (C/N).

The equation below calculates the C/N of an amplifier:

$$C/N \text{ (dB)} = C_i \text{ (dBmV)} + 57.4 - NF \text{ (dB)} \quad (1)$$

Where:

- $C_i$ : input signal

- *NF*: Noise figure of the amplifier

Note: the number 57.4 is the noise power for QAM carriers. This value will vary marginally depending on the temperature.

Equation (1) shows the significance of the Rx power versus noise figure of the amplifier, in overall system design.

The overall system C/N can be derived from the following equation:

$$C/N_{total} (dB) = -10 \log \left\{ 10^{\frac{-C/N_1}{10}} + 10^{\frac{-C/N_2}{10}} + \dots + 10^{\frac{-C/N_n}{10}} \right\} \quad (2)$$

Where  $C/N_x$  is the carrier to noise of each amplifier calculated independently.

When cascading identical amplifiers, the following approximation is typically used:

$$C/N_{total} (dB) = C/N_x - 10 \log n \quad (3)$$

Where:

- $C/N_x$ : the carrier to noise of a single amplifier
- $n$ : the number of identical amplifiers in cascade.

### **2.4.3. Distortion**

The build-up of distortions in a cascaded plant are less predictable than noise. The following equation can be used to estimate the carrier to composite triple beat (C/CTB) and carrier to composite second order distortion (C/CSO):

$$C/CTB_{total} (dB) = C/CTB_x - 20 \log n \quad (4)$$

$$C/CSO_{total} (dB) = C/CSO_x - 10 \log n \quad (5)$$

Where:

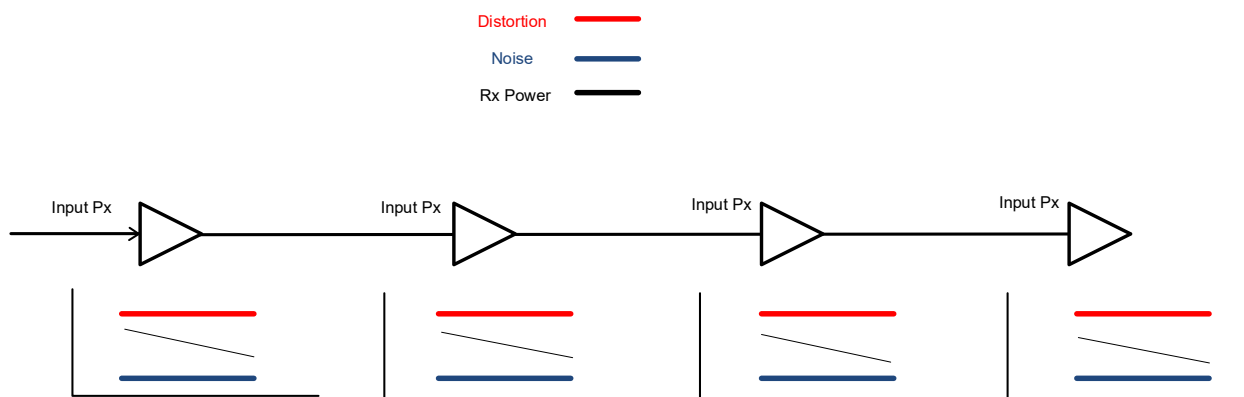
- $C/CTB_x$ : the distortion of a single amplifier
- $C/CSO_x$ : the distortion of a single amplifier
- $n$ : number of identical amplifiers in cascade

### **2.4.4. Designing a Noise-Limited System**

For optimal performance, operators design systems that are unity gain. This means that the loss between two amplifiers is equal to the gain of each amplifier. If the loss is less than the gain, the distortions will accumulate. Whereas if the loss is greater than the gain then the input power will be less than the desired amount, degrading the system C/N.

Due to the difficulties that come with designing a system that is both noise and distortion limited, removing one of those parameters will be optimal. Given that noise performance of amplifiers is far more straight-forward in comparison to distortion, designing a noise-limited system is a very attractive idea.


Since distortions are highly dependent on output power, designing a noise-limited system can be achieved by reducing the output power out of the node/amplifiers and making sure the signal level received at the next amplifier is high enough, based on equation 1. Figure 5 demonstrates balancing the signal level against noise and distortion in a system design.



**Figure 5 – Signal Level Balanced Between Noise and Distortion**


## 2.5. Typical Plant Topologies

In this paper we focus on Node + 0 (Figure 6) and Node + 4 (Figure 7) networks as the basis for performance evaluation.



**Figure 6 – Node + 0 Passive HFC Plant Topology**





**Figure 7 – Node + 4 Cascade Plant Topology**

This plant topology has 5 amps total in cascade, the node amp and the 4 network amps. The network is built with a cascade of 4 trunk spans and single distribution span as described in section 2.4.

## 2.6. Problem Definition

As described before, the ESD communication system is limited in its transmit power. Optimal allocation of available transmit power and appropriate bit-loading (profile definition) is needed to get the maximum data rate out of the network. In this paper, we outline a framework for closed loop optimization of the throughput of ESD system subjected to TCP constraint mentioned above. We will show that combining careful allocation of channels, closed loop optimization of transmit power, and adaptive bit-loading achieves considerable gains in data rate and reduction in TCP for network topologies currently present in MSO networks.

## 3. Theoretical Framework for Closed Loop Throughput Optimization

In this section, we present the theoretical framework for optimizing the available transmit power for throughput of the cable network. We start off with the simpler case of Node + 0 passive network and then extend the theory to cover general case of Node + X, X>0, networks.

### 3.1. Node + 0 Network

The capacity evaluation for the extended spectrum HFC network requires knowledge of the channel characteristics and capacity limiting factors. The capacity limiting factor in the transmitter is amplifier distortion. At the receiver, additive white Gaussian noise and receiver distortion due to analog-to-digital conversion limits capacity. HFC transmission schemes such as DOCSIS 4.0 [4] use OFDM modulation, where the channel is partitioned into K narrowband subcarriers  $k = 1, \dots, K$  with a subcarrier spacing  $\Delta f$ . Those orthogonal channels are coupled only by nonlinear distortion or a sum power constraint. The transmit power per carrier  $x(k)$  as well as the information rate per carrier  $b(k)$  can be adjusted per carrier. The data rates for a given signal-to-noise ratio  $SNR(k)$  on carrier k is given by

$$R = \eta \Delta f \sum_{k=1}^K \min \left( \log_2 \left( 1 + \frac{SNR^{(k)}}{\Gamma} \right), b_{\max} \right). \quad (6)$$

where limitations of modulation and coding are considered in terms of an SNR gap to capacity  $\Gamma$  [12] as well as with a limit  $b_{\max}$  to the number of bits transmitted per carrier and channel use. The OFDM system requires overhead for the cyclic extension to guarantee orthogonal channels, which is considered

in an efficiency factor  $\eta$ . Using  $\eta = 1$ ,  $\Gamma = 1$  and  $b_{max} \rightarrow \infty$  gives the capacity without coding and modulation limitations.

Capacity, C, and achievable rate, R, are evaluated with respect to power constraints where the simplest case is a sum power constraint [13]. For practical systems, additional per-carrier constraints are considered, as shown in [14], Sec. 3.1.6. For this case, achievable rate, R, and capacity C for the case of  $\Gamma = 1$ , is the solution to

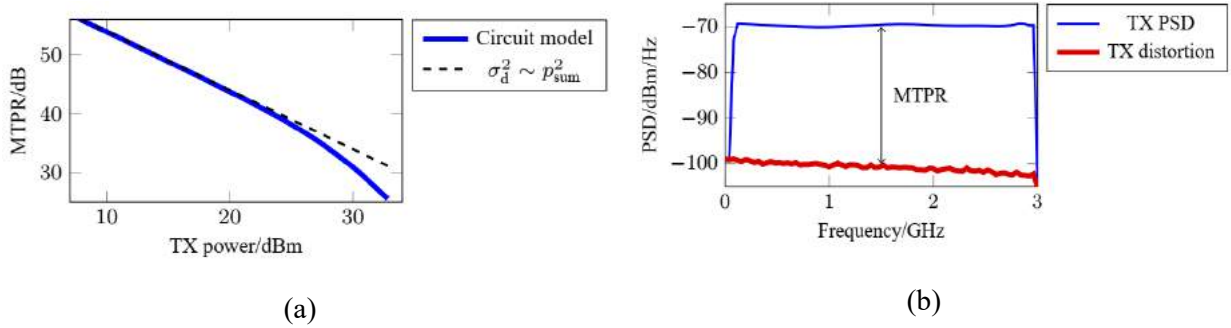
$$\begin{aligned} R_l = \max_{x_l^{(k)}} \sum_k \log_2 \left( 1 + \frac{|H_l^{(k)}|^2 x^{(k)}}{\Gamma \sigma_l^{(k),2}} \right) \\ \text{s.t. } \sum_k x^{(k)} \leq p_{max} \\ \text{s.t. } 0 \leq x^{(k)} \leq p_{mask}^{(k)} \end{aligned} \quad (7)$$

where  $H_l^{(k)}$  is the channel coefficient on carrier k (attenuation, phase) between node and cable modem (CM) l and  $\sigma_l^{(k),2}$  is the additive white Gaussian noise variance on carrier k for CM l.

The power constraints are formulated as a sum power limit  $p_{max}$  and a spectral mask constraint  $p_{mask}^{(k)}$ . Limitations of the modulation alphabet size to  $b_{max}$  are incorporated into the spectral mask constraint, using  $p_{mask}^{(k)} = \Gamma(2^{b_{max}} - 1)\sigma_l^{(k),2} / |H_l^{(k)}|^2$ . The solution to Eq. (7) is obtained by a modified water-filling algorithm as described in [14], chapter 3.1.6. Other algorithms to solve Eq. (7) have been published in [15],[17].

The sum power limit,  $\sum x^{(k)} \leq p_{max}$ , in Eq. (7) can be seen as a simplified model for the behavior of real transmit amplifier, where the SNR and thus the data rate is limited by distortion increasing with increasing transmit power. Transmit amplifier distortion can be seen as a transmit power dependent noise source with variance  $\sigma_d^2$ . It is characterized in measurement and simulation by a missing tone power ratio (MTPR). MTPR is the ratio between signal power and distortion power, as shown in Figure 8 (b). It is determined as the signal level on one OFDM carrier which is transmitted with zero power while the others are transmitted at the desired level ( $MTPR^{(k)} = x^{(k)} / \sigma_d^2$ ). In the following discussion, the cable modem index l is skipped without loss of generality.

Figure 8 (a) shows the increase of nonlinear distortion  $\sigma_d^2$  in a 3 GHz amplifier circuit model. The MTPR decreases with increasing transmit power,  $p_{sum} = \sum_{k=1}^K x^{(k)}$ . In frequency domain, as shown in Figure 8 –(b), distortion is approximately flat.



**Figure 8 – Distortion evaluation from amplifier circuit model**

The dependency between distortion variance  $\sigma_d^2$  is and signal power  $p_{\text{sum}}(x^{(k)})$  can be described by  $\sigma_d^2 = \delta \left( p_{\text{sum}}(x^{(k)}) \right)^\alpha$ . For the amplifier shown in Figure 8 (a), the constants are  $\delta = -64$  dB and  $\alpha = 2$ . Following the argumentation of [9] a lower bound for the capacity of the nonlinear copper channel is derived.

Introducing distortion in the SNR per carrier gives the term

$$SNR^{(k)} = \frac{|H^{(k)}|^2 x^{(k)}}{\sigma^2 + \delta^{(k)} (p_{\text{sum}}(x^{(k)}))^\alpha} \quad (8)$$

where the distortion variance is  $\sigma_d^{(k),2} = \delta^{(k)} \left( p_{\text{sum}}(x^{(k)}) \right)^\alpha$ , assuming white distortion. This gives the rate  $R$  (or capacity  $C$  with  $\Gamma = 1$ ) according to

$$R = \max_{x^{(k)}} \sum_k \log_2 \left( 1 + \frac{|H^{(k)}|^2 x^{(k)}}{\Gamma (\sigma^2 + \delta^{(k)} (p_{\text{sum}}(x^{(k)}))^\alpha)} \right) \\ \text{s.t. } 0 \leq x^{(k)} \leq p_{\text{mask}}^{(k)} \quad (9)$$

The derivative  $\frac{\partial R}{\partial x^{(k)}}$  is given by

$$\frac{\partial R}{\partial x^{(k)}} = \frac{|H^{(k)}|^2}{(\sigma^2 + \delta_k (p_{\text{sum}}(x^{(k)}))^\alpha) (\Gamma + SNR_k)} - \\ \sum_{d=1}^K \frac{|H^{(d)}|^2 x^{(d)} \alpha \delta^{(d)} (p_{\text{sum}}(x^{(k)}))^{\alpha-1}}{(\sigma^2 + \delta^{(d)} (p_{\text{sum}}(x^{(k)}))^\alpha)^2 (\Gamma + SNR^{(d)})} \quad (10)$$

and  $\partial R / \partial x^{(k)} = 0$  must hold for the optimal power allocation for all carriers with  $0 < x^{(k)} < p_{\text{mask}}^{(k)}$ . The optimal power allocation can be found, e.g., by a projected gradient method with a step size  $\rho$  as given by

$$x_{t+1}^{(k)} = \min \left( \max \left( x_t^{(k)} + \rho \frac{\partial R}{\partial x^{(k)}}, 0 \right), p_{\text{mask}}^{(k)} \right) \quad (11)$$

It can be shown [17] that the solution to equation (10) takes the following form,

$$\frac{1}{\mu} = \sum_{d=1}^K \frac{|H^{(d)}|^2 x^{(d)} \alpha \delta^{(d)} (p_{\text{sum}}(x^{(k)}))^{\alpha-1}}{(\sigma^2 + \delta^{(d)} (p_{\text{sum}}(x^{(k)}))^{\alpha})^2 (\Gamma + \text{SNR}^{(d)})} \quad (12)$$

The dependency between  $\mu$  and  $p_{\text{sum}}$  is given by,

$$\frac{1}{\mu} = \frac{1}{|I_{\text{fill}}|} \left( p_{\text{sum}} + \sum_{k \in I_{\text{fill}}} \frac{\Gamma \sigma_{nd}^{(k),2}}{|H^{(k)}|^2} - \sum_{k \in I_{\text{mask}}} p_{\text{mask}}^{(k)} \right) \quad (13)$$

Where  $I_{\text{mask}}$  is the set of subcarriers for spectral mask constraint is active and  $I_{\text{fill}}$  is the set of subcarriers with power allocation to meet water-fill level, and  $|I_{\text{fill}}|$  denotes cardinality (number of elements) of the set  $|I_{\text{fill}}|$ .

Hence the transmit power per subcarrier is given by,

$$x^{(k)} = \begin{cases} \frac{1}{\mu} - \frac{\Gamma \sigma_{nd}^{(k),2}}{|H^{(d)}|^2} & \text{for } k \in I_{\text{fill}} \\ 0 & \text{for } k \in I_0 \\ p_{\text{mask}}^{(k)} & \text{otherwise} \end{cases} \quad (14)$$

Where  $I_0$  is the set of subcarriers where positiveness constraint given in equation (7) is active and therefore no power is allocated.


To implement the optimization scheme, the distortion parameters,  $\delta$  and  $\alpha$  must be known from an amplifier characterization. During operation, the noise conditions must be known from an SNR measurement. The algorithm performs multiple water-filling steps. As the optimization can be done by software in background during operation of the link, there is no issue with computation time.

### 3.2. Extention to Node + X, X > 0, Networks

The algorithm described in section 3.1 assumes a single source of transmitter distortion at the node and a single receiver noise source at each CM, as it is reflected in the SNR in Eq. (8). In case of a Node + X topology with multiple intermediate amplifiers, each intermediate amplifier represents an additional source of distortion and each amplifier input experiences additive receiver noise.

Still, it can be shown that under a certain assumption, the optimization framework of section 3.1 can be applied to the Node + X case, too. The precondition for the spectrum optimization to be applied is that the attenuation and down-tilt of the cable section, given by the transfer function  $H_{\text{span}}^{(k)}$ , is compensated by

amplifier stage  $H_{amp}^{(k)}$  at the end of the cable section, such that the power level and spectral shape of the transmit signal is the same at each intermediate amplifier  $H_{span}^{(k)} H_{amp}^{(k)} = 1$ , as shown in Figure 9.



**Figure 9 – Node+X topology with Power Spectrum Optimization**

Mathematically, the transmit power per carrier  $k$ ,  $x^{(k)}$  is (approximately) the same at node output and the amplifier outputs,  $x^{(k)} \approx x_{amp,1}^{(k)} \approx \dots \approx x_{amp,n}^{(k)} \forall k = 1, \dots, K$ . With this condition satisfied, all the amplifier distortion of  $N$  amplifiers can be combined into one distortion  $(N + 1) \delta^{(k)}(p_{sum}(x^{(k)}))$ .

The receiver noise of the amplifiers, assuming  $\sigma_{amp}^2 \ll |H_{span}^{(k)}|^2 x^{(k)}$  for all the used frequencies, can be handled as additive noise. The receiver noise of the amplifiers is summed up, referred to the CM receiver and added to the CM receiver noise  $\sigma^2$  which gives the overall additive noise term to be

$$N \sigma_{amp}^2 |H_{amp}^{(k)}|^2 |H^{(k)}|^2 + \sigma^2.$$

Accordingly, the algorithm of section 3.1 remains as is while performing the following substitutions

$$\text{Receiver noise: } \sigma^2 \quad \rightarrow \quad N \sigma_{amp}^2 |H_{amp}^{(k)}|^2 |H^{(k)}|^2 + \sigma^2$$

$$\text{Transmitter distortion: } \delta^{(k)}(p_{sum}(x^{(k)})) \quad \rightarrow \quad (N + 1) \delta^{(k)}(p_{sum}(x^{(k)}))$$

Compared to the Node + 0 architecture, transmitter distortion optimization is even more relevant in the Node + X architecture, as there is a higher distortion level present due to the distortion of multiple amplifiers adding up. It may also be beneficial to drive the TCP of individual amps down as much as possible to lower individual nonlinear distortion contributions from amps. We will look at strategies on achieving this objective in this in this paper.

### 3.3. Algorithm Description

The above framework leads to an iterative throughput optimization algorithm described below.

**Initialization Step:** Initialize the TCP,  $P_{sum}(0)$ , to the target maximum power level. Index 0 refers to initial value

Following that, the following steps are performed in an iterative loop, until the  $P_{sum}(n)$  converges to a steady value.  $|P_{sum}(n) - P_{sum}(n + 1)| < Threshold$

**Iterative Step 1:** Apply water-filling described in [14] with the current  $P_{sum}$  to discover,

- $I_0$  = The set of subcarrier that get no power allocation
- $I_{fill}$  = The set of subcarriers that gets power allocation to water-fill level
- $I_{mask}$  = The set of subcarrier that gets power allocated to  $p_{mask}^{(k)}$  level. For these subcarrier  $p_{mask}^{(k)}$  level is hit before reaching water-fill level and therefore no additional power is wasted to reach water-filling level

**Iterative Step 2:** Based on subcarrier sets information from previous step and Equation (14), calculate the power allocation per subcarrier,  $x^{(k)}$ .

**Iterative Step 3:** Calculate water-filling level,  $\frac{1}{\mu}$ , using equation (12)

**Iterative Step 4:** Update the TCP,  $P_{sum}(n + 1)$ , from equation (13)

**Exit Criteria:** If  $|P_{sum}(n) - P_{sum}(n + 1)| < Threshold$  then stop iteration. Otherwise go back to Iterative Step 1.

Pseudo code for the algorithm is given below.

**Table 1 – Optimization algorithm summary**

|                |                                                                                                                                                                                                                                                                                                                                       |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Initialization | $P_{sum}(0) = \text{max allowed TCP}$                                                                                                                                                                                                                                                                                                 |
| Iteration      | <p>Identify the sets <math>I_0</math>, <math>I_{fill}</math>, and <math>I_{mask}</math> (water-filling, [14])</p> <p>Update <math>x^{(k)}</math> using Eq. (14)</p> <p>Calculate <math>\mu</math> from Eq. (12) with updated <math>x^{(k)}</math></p> <p>Update <math>P_{sum}(n + 1)</math> from <math>\mu</math>, using Eq. (13)</p> |
| Exit Criteria  | $ P_{sum}(n) - P_{sum}(n + 1)  < Threshold$                                                                                                                                                                                                                                                                                           |

## 4. System Level Solution for Closed Loop Optimization

In order to apply the closed loop optimization algorithm developed above in a practical system, we first need to consider the impact of channel allocation. Channel allocation and power optimization together form our overall closed loop solution. We consider two power allocation strategies; Full spectrum allocated to each CM, and channel stacking or staggered channel allocation.

### 4.1. Full Spectrum allocated to each CM

Consider the mid-split US/DS partitioning with all the CMs allowed to use DS OFDM channels anywhere in the spectrum from 108 MHz to 1794 MHz. In this case, power allocation is optimized considering channel frequency response and noise of all channels for a CM connected to a particular tap. Given that each channel is potentially shared between all CMs in the node, we have to carefully select the CM that we target the optimization algorithm for (i.e. what channel frequency response and noise responses to use in the algorithm). Going for the worst-case CM (farther away from node) or the best-case CM (closer to node) may not lead to overall node throughput optimization. A CM that represents median or overage behavior, in terms of received signal quality, would lead to better results. Once the downstream power is optimized, DOCSIS has other tools, such as the profiles, to fine tune the throughput for CMs with different received signal quality.

### 4.2. Channel Stacking

In this case, based on the observation that the cable channel has higher losses at higher frequencies, we allocate the lower frequency channels to far away CMs and higher frequency channels to close in CMs. It will be shown later that this allows us to reduce the TCP of node and amps significantly, opening up potential other benefits in overall network architecture.

### 4.3. Implementation Considerations

Closed loop algorithm implementation considerations are described in this section. The aim is, as much as possible, to work within the current DOCSIS 4.0 standard provisions to implement the algorithm described in previous section. We also highlight aspects of the standard that can be improved to better facilitate the closed loop throughput optimization.

The algorithm described in section 3.3 requires per subcarrier channel frequency response estimates and total noise estimates from the CMs to derive the optimal power allocation. DOCSIS 3.1 and 4.0 provide following proactive network maintenance (PNM) features to help gather this information from the CM:


- Downstream Channel Estimate Coefficients: CMTS can command CM to send Channel Estimate coefficients to CMTS (section 9.3.4 of [3], [4]).
- Downstream Receive Modulation Error Ratio (RxMER) Per Subcarrier: CMTS can command CM to send MER estimates to CM (section 9.3.6 of [3], [4])

Noise power per subcarrier can also be derived from the above information.

Once the optimal power allocation solution is found, the CMTS needs to shape the downstream transmit spectrum accordingly. The DOCSIS 4.0 spec [4] has provisions for the node to shape the transmit RF spectrum under the following constraints:

- Apply a uniform up tilt (i.e. same gradient across all entire spectrum) to the transmit spectrum so that more power is allocated to higher frequencies.
- Introduce multiple step downs/ups to maintain total composite power bound (sum step downs < 10 dB)
- Spectrum discontinuities, step down/up, are only allowed in OFDM channel boundaries

Figure 10 shows an example power allocation instance allowed by the spec viewed at interface C – before the uniform tilt and power amplification (PA).



**Figure 10 – Downstream Transmit Power allocation as seen from Node Interface C**

At interface D (node output port), post tilt application and PA, the above spectrum appears with a uniform up tilt while keeping the same step-downs seen in Figure 10.

One key area where the spec could be amended is to allow for more flexible spectrum shapes with possible flat power spectral density (PSD) at high frequencies at interface D. This means removing the condition of uniform tilt across the entire spectrum mentioned above. This will allow for more accurate implementation of some of the transmit PSD optimization scenarios described later in this document. However, we should weigh the benefits of doing so against potential complications to the node and amp architectures as well as the operation of the network.

We can bring in a machine learning approaches to incorporate various other aspects of the system in the overall solution, such as:

- Take into account the individual CM throughput usage over time in channel allocation
- Overall node-wide throughput usage over time
- Service agreement data for individual CMs
- Prior knowledge of network topology

A profile management application (PMA) can also be used in conjunction with the items mentioned above to increase the overall performance in the distribution plant. By enabling dynamic bit-loading and profiles assigned to each service group, the overall plant throughput and stability of the plant will increase.

#### **4.4. Create Headroom in Power Budget for Soft Flexible MAC Architecture**

In this section, we briefly address an added benefit of potential TCP reduction by closed loop optimization in creating room in the overall network power budget for a soft flexible MAC architecture (FMA) solutions. Currently the node power budget is very tight already with remote PHY device (RPD) solutions. With a node power budget ranging from 160 W to 180 W for North America, the RPD and the RF power amplifiers are already using up most of this power. For example, for 85% power delivery efficiency, power left over for other potential uses are roughly 16 W and 33 W for a 2x2 node with 4 legs for overall 160 W and 180 W power budgets respectively. A large chunk of the power is taken up by the 4 PAs, which are assumed to be at 71 dBmV TCP at the node output port.

For FMA solutions, especially remote MAC device (RMD), where MAC functionality is also distributed to the node in addition to PHY, the above excess power is not sufficient for a fully software based solution using general purpose compute, let alone having additional headroom for future edge computing applications. This forces a solution which is either based on ASIC/FPGA or/and a CPU with limited compute (to fit within available power budget).



If we can lower the TCP by just 2 dB, to 69 dBmV, this increases the available power for FMA and edge compute to 28 W to 45 W. As shown in section 5, with the additional power saving achievable with closed loop optimization combined with channel stacking, it is possible to reduce the TCP by as much as 6 dB for 1.8 GHz ESD. Even for 1.2 GHz DOCSIS, we can explore channel stacking with power optimization to get some reduction in power consumption. This opens the possibility of a soft MAC architecture and also leaves enough headroom for other future edge compute applications.

## 5. Simulation Results and Discussion

Rate results for different network topologies, channel allocation strategies and power allocation strategies are given in this section. Node + 0 and Node + 4 network topologies are considered. Allowing all CMs to access the entire downstream spectrum vs channel stacking is also considered. In terms of power allocation, three different allocation strategies are compared:

- Tilted TX PSD to receive flat RX spectrum at CM.
- Flat TX PSD as implied by conventional water filling solution.
- Optimal power allocation based on the algorithm described in this paper.

We explore both the mid and high split plant scenarios:


- Mid-split: DS starts at 108 MHz
- High-split: DS starts at 258 MHz

### 5.1. Full Optimization for Node + 0 and Node + 4

In this section, we explore the optimization of transmit PSD for capacity without any other constraints imposed by operational considerations. This is partly an academic exercise to gain insight into the properties of the power allocation algorithm

#### 5.1.1. Node + 0 Network, Mid-Split

Figure 11 shows the throughput for CMs connected to different taps for the three different power allocation strategies.




**Figure 11 – Rate for flat RX (blue), flat TX (red) and optimal (yellow) power allocation**


Throughput gain from using the optimal algorithm is limited in this case.

### 5.1.2. Node + 4 Network, Mid-Split

The same three power allocation strategies described in section 5 are used here. It is assumed that the distribution cable sections have 35 dB attenuation at 1 GHz. The Node + 4 network is constructed with four straight trunk cable sections with amplifiers followed by a cable section.




**Figure 12 – Attenuation of the cable sections (left) and signal and noise PSDs at the last Tap (right)**



**Figure 13 – Node + 4 mid split SNR at Tap 3 (left) and bit allocation at Tap 3 (right)**


The distortion of all 5 transmit amplifiers (CMTS and 4 amplifiers) is the same. Each amplifier compensates the channel attenuation and tilt of preceding cable segment completely (segment unity gain) such that the transmit spectrum is the same at each stage. For the noise figure of the amplifiers, 2 amplifiers with 5 dB and 2 amplifiers with 10 dB are assumed. Additional 3 dB of losses in the receiver is assumed (account for losses in passive connectors, diplexers, etc).

Figure 14 shows the throughput for CMs connected to different taps for the 3 power allocation strategies. Up to 10% Rate improvement is achieved with the optimal power allocation.



**Figure 14 – Node + 4 mid-split Data rates for tilted TX- flat RX PSD (blue), flat TX PSD (red), and optimized PSD (yellow)**

Figure 15 shows transmit power spectrum for the 3 power allocation schemes (left) and TCP when optimizing throughput for a certain tap (right). In case of tilted TX PSD and flat TX PSD, the transmit power spectrum and TCP don't depend on the tap. When optimizing throughput for long taps (5, 6), the optimal spectrum shape is flat, while optimizing for short taps (1) gives a tilt for all frequencies as an optimal shape. When optimizing for the medium taps (2-4), the optimal spectrum shape has a tilt at low frequencies followed by flat region at high frequencies. The tilt at low frequencies indicates that the performance at these frequencies are nonlinear distortion dominated. On the other hand, at high frequencies the dominance of thermal noise (i.e. signal is more attenuated and hence closer to thermal noise floor) makes the flat spectrum optimal (classic water-filling comes into play). TCP graphs shows the TCP at node when optimized for tap number given in x-axis.

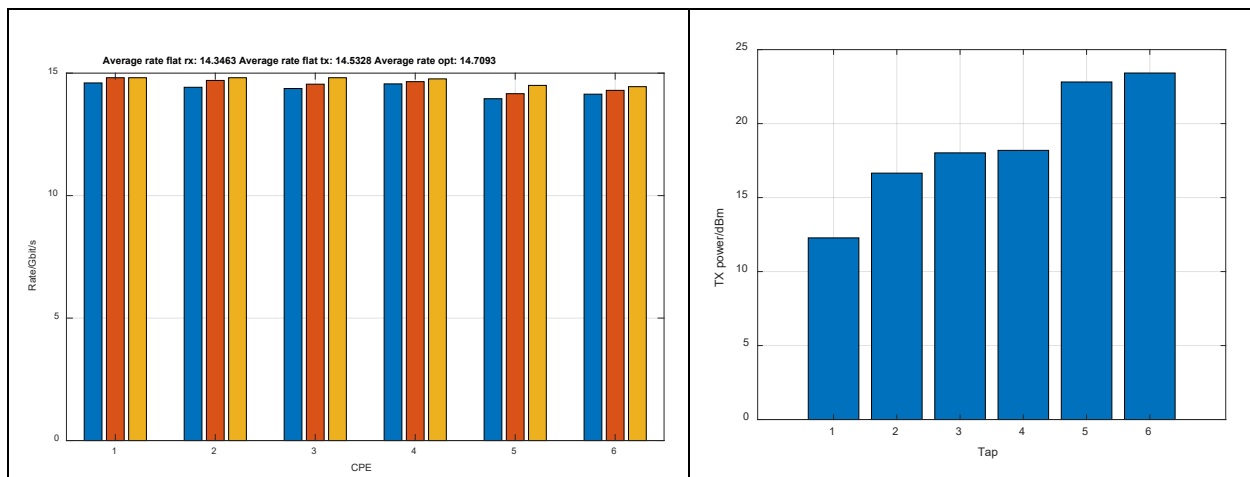


**Figure 15 – Node + 4 mid-split Transmit PSDs and TCPs for different allocation schemes**

TCP requirement is dominated by the far away CMs, requiring the full 71 dBmV (22.25 dBm) of power.


### 5.1.3. Node + 0 Network, High-Split

We've repeated the test for high split note to get the following rate and TCP results. TCP requirement is dominated by far away CMs demands.



**Figure 16 – Node + 0 high-split Data rates for flat RX (blue), flat TX (red), and optimized (yellow) and corresponding TCP**

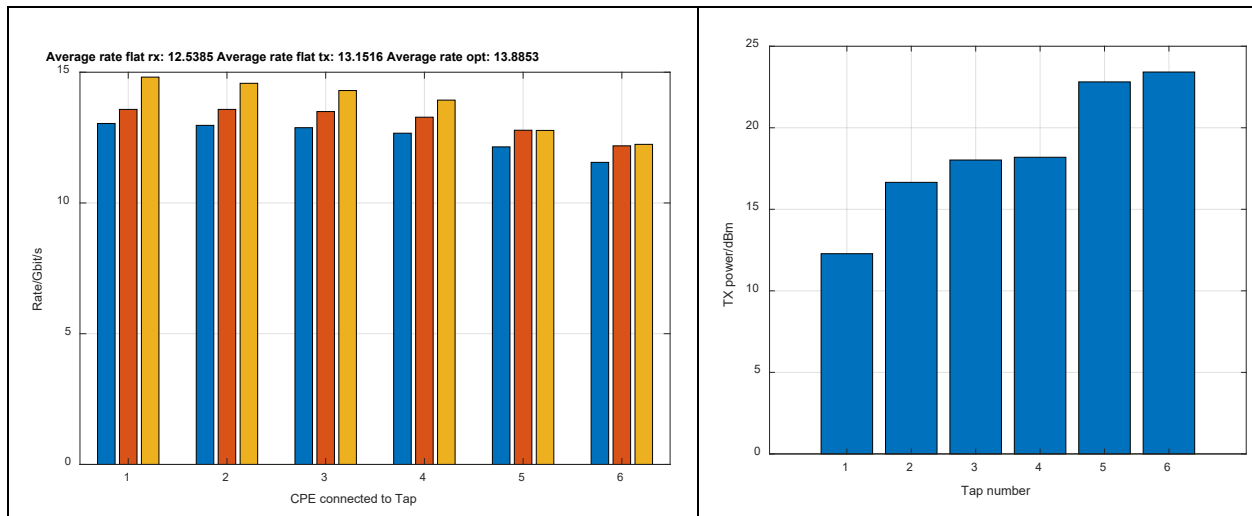
Optimal PSDs are shown in Figure 17. Although we can maintain a uniform tilt in spectrum for the close in and medium range CMs, long range CMs forces a flat spectrum at high frequencies.



**Figure 17 – Node + 0 high-split Transmit PSDs for different allocation schemes**


#### **5.1.4. Node + 4 Network High-Split**

Up to 10% rate improvement achieved with optimal power allocation as show in Figure 18 (left).



**Figure 18 – Node + 4 high-split Data rates for flat RX (blue), flat TX (red) and optimized (yellow) and the corresponding TCP**

Power allocation across frequency for the 3 allocation schemes is shown in Figure 19.



**Figure 19 – Node + 4 high-split Transmit PSDs for different allocation schemes**

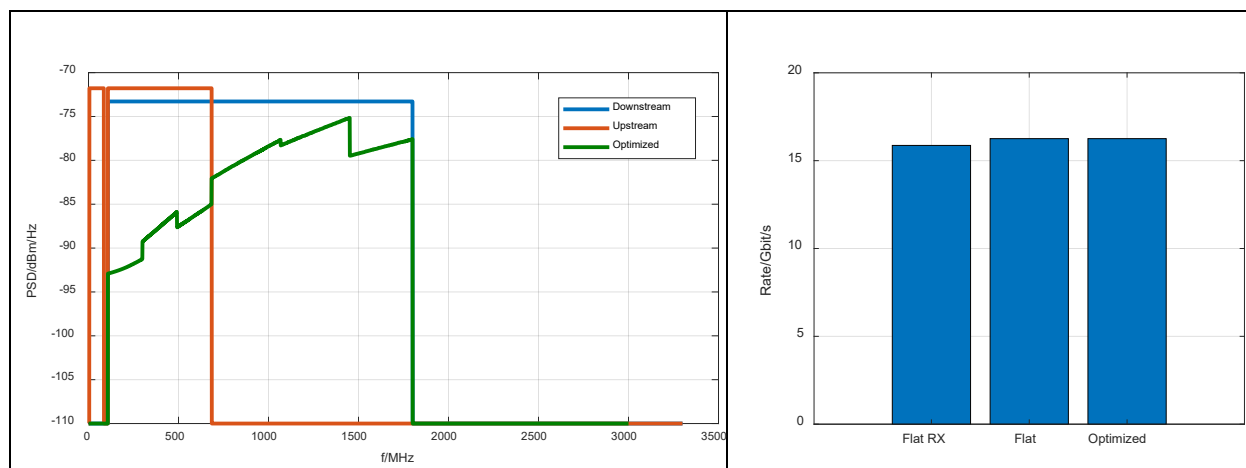
## 5.2. With Staggered Channel Allocation – Channel Stacking

In these tests, we exploit the property of the cable channel that it has higher losses at higher frequencies. To minimize the impact of higher losses in higher frequencies, we allocate these higher frequency channels to close in CMs and lower frequency channels to far away CMs. Consider to hypothetical channel allocation (non-DOCSIS) shown in Figure 20. The channel estimate for each allocated channel is used in the optimization algorithm.

| Modem<br>Channel | 1-4 (Tap 1) | 5-8 (Tap 2) | 9-12 (Tap 3) | 13-16 (Tap 4) | 17-20 (Tap 5) | 21-24 (Tap 6) |
|------------------|-------------|-------------|--------------|---------------|---------------|---------------|
| 1                |             |             |              |               |               |               |
| 2                |             |             |              |               |               |               |
| 3                |             |             |              |               |               |               |
| 4                |             |             |              |               |               |               |
| 5                |             |             |              |               |               |               |
| 6                |             |             |              |               |               |               |
| 7                |             |             |              |               |               |               |
| 8                |             |             |              |               |               |               |
| 9                |             |             |              |               |               |               |

**Figure 20 – Channel to CM allocation for PSD optimization**

Figure 21 shows PSD for the above channel allocation scheme. A significant reduction in TCP (up to 6 dB) is achieved through this channel allocation scheme without losing any throughput.



**Figure 21 – Power allocation vs Data rates**

However, this is not quite a practical allocation scheme. In practice we would allocate set of channels to a group of CMs and take advantage of statistical multiplexing to improve the utility of the spectrum.

Figure 22 shows a practical channel allocation scheme where channels allocated to long range CMs are indicated in green and channels allocated to short range CMs are indicated in yellow. Within each channel set, the CM corresponding to the channel estimate used in the optimization algorithm is indicated in dashed lines.

| Modem<br>Channel | 1-4 (Tap 1) | 5-8 (Tap 2) | 9-12 (Tap 3) | 13-16 (Tap 4) | 17-20 (Tap 5) | 21-24 (Tap 6) |
|------------------|-------------|-------------|--------------|---------------|---------------|---------------|
| 1                |             |             |              |               |               |               |
| 2                |             |             |              |               |               |               |
| 3                |             |             |              |               |               |               |
| 4                |             |             |              |               |               |               |
| 5                |             |             |              |               |               |               |
| 6                |             |             |              |               |               |               |
| 7                |             |             |              |               |               |               |
| 8                |             |             |              |               |               |               |
| 9                |             |             |              |               |               |               |

**Figure 22 – Channel stacking for two service groups**

Note that each CM is allocated 5 OFDM channels, which will allow us to maintain 10 Gbit/s for the service group (SG).

SG A:

- CMs in Taps 1 to 3.
- Allocated channels 5 to 9

SG B:

- CMs in Taps 4 to 6
- Allocated channels 1 to 5

This channel allocation enables what we call 10G DOCSIS and it will appear like a node-split in the sense that we are supporting two 10G SGs using available spectrum.

By optimizing the power allocation for each channel to match at least one of the taps in each SG, we are optimizing the average throughput of each SG (assuming each user has roughly the same probability of using each channel).

The alternative scheme shown in Figure 23 uses the same channel allocation, but the channel frequency responses used in the optimization algorithm is taken from the medium range CM within each SG. This is the default channel allocation method used in following tests unless mentioned otherwise.

| Modem<br>Channel | 1-4 (Tap 1) | 5-8 (Tap 2) | 9-12 (Tap 3) | 13-16 (Tap 4) | 17-20 (Tap 5) | 21-24 (Tap 6) |
|------------------|-------------|-------------|--------------|---------------|---------------|---------------|
| 1                |             |             |              |               |               |               |
| 2                |             |             |              |               |               |               |
| 3                |             |             |              |               |               |               |
| 4                |             |             |              |               |               |               |
| 5                |             |             |              |               |               |               |
| 6                |             |             |              |               |               |               |
| 7                |             |             |              |               |               |               |
| 8                |             |             |              |               |               |               |
| 9                |             |             |              |               |               |               |

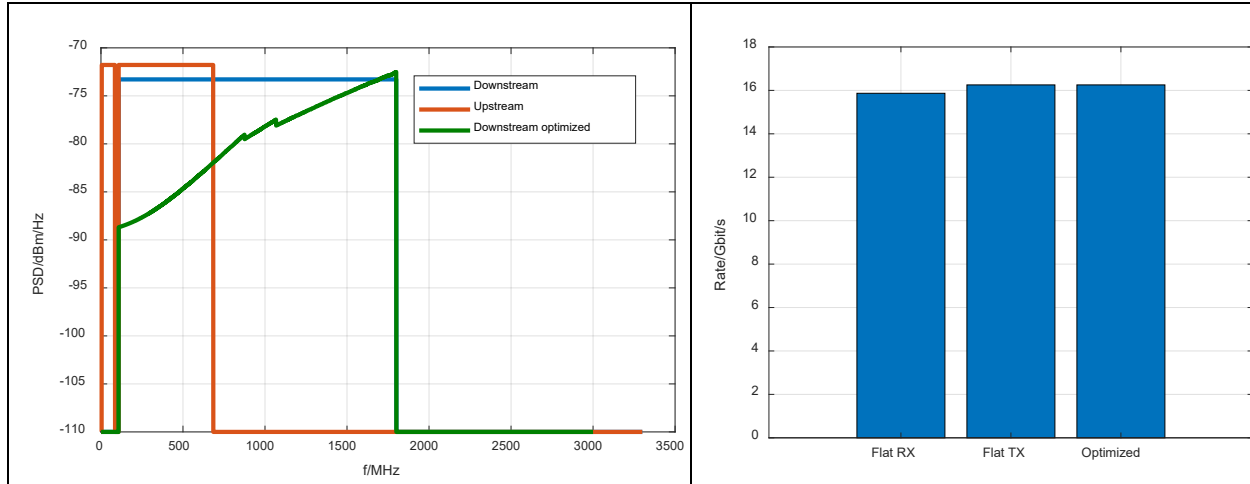
**Figure 23 – Default channel stacking scheme for two service groups**

In practice, with additional information, we could make more data driven approach to allocating channels and deciding which tap/CM channel estimate we use for optimization.



### 5.2.1. Node + 0 Network Mid-Split

Figure 24 shows the optimal PSD and resulting data rates for the channel allocation scheme given in Figure 23 for Node + 0 Network.

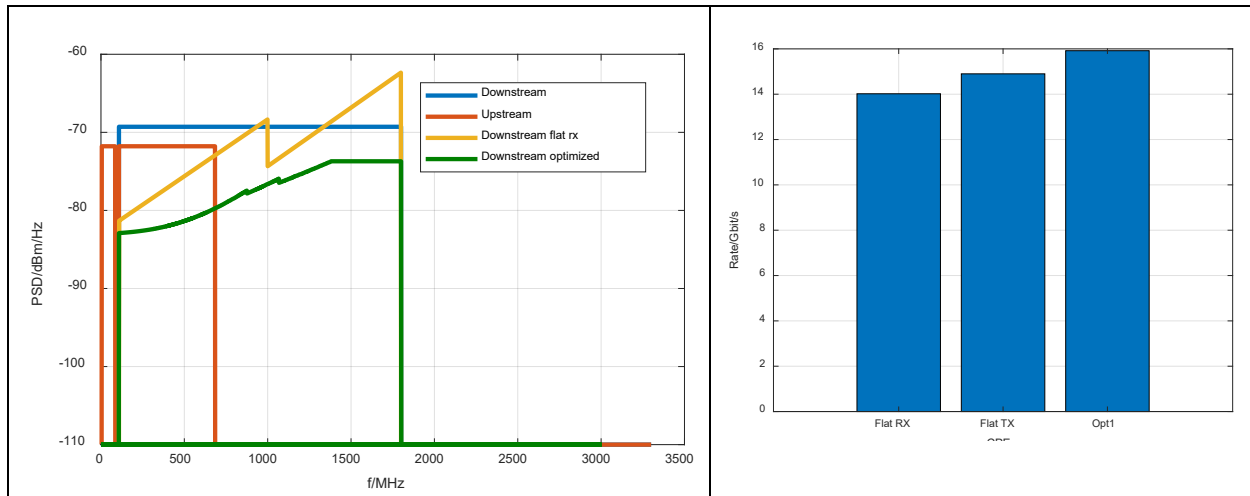


**Figure 24 – Node + 0 mid-split Power allocation and Rate for channel stacking**

Note that the TCP has reduced by nearly 6 dB compared to assigning all the channels to each CM. Furthermore, the overall rate is higher compared to allocating all channels to each CM.

### 5.2.2. Node + 4 Network Mid-Split

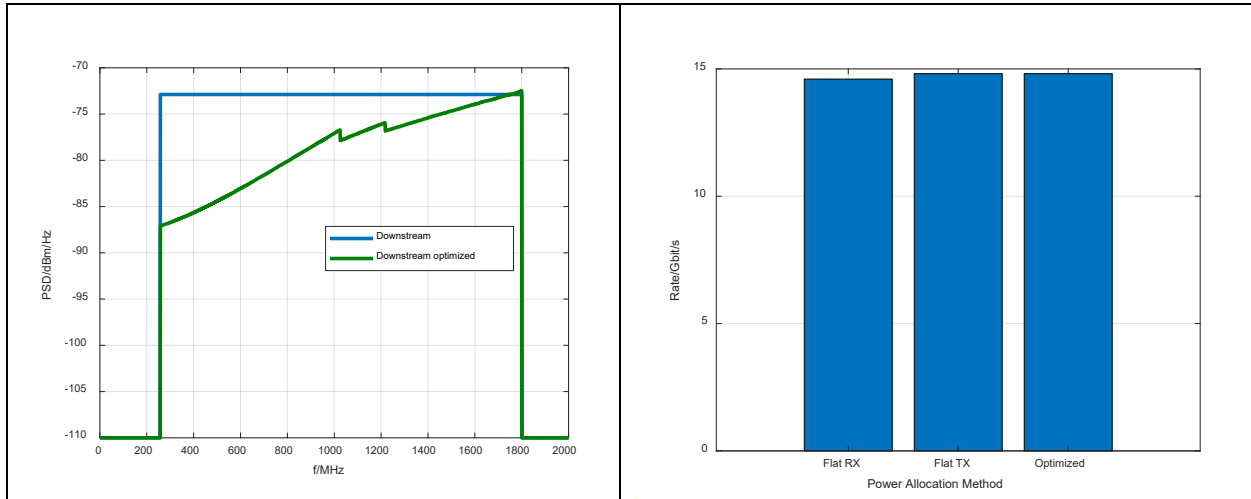
Channel allocation based power optimization results for Node + 4 mid-split case is shown in Figure 25. As in Node + 0 case, we are making significant savings in TCP. In addition, the optimal power allocation gives 10-15% increase in the throughput compared to non-optimal power allocation schemes.



**Figure 25 – Node + 4 mid-split power allocation and Rate for channel stacking**

### 5.2.3. Node + 0 Network High-Split

Channel allocation based power optimization results for Node + 0 high-split case is shown in Figure 26. As with Node + 0 mid-split cases, we are making significant savings in TCP while not losing any throughput.




**Figure 26 – Node + 0 high-split power allocation and Rate for channel stacking**

The optimized TX TCP is 6 dB below the TX power used for flat TX and flat RX.

### 5.2.4. Node + 4 Network High-Split

Channel allocation based power optimization results for Node + 4 high-split case is shown in Figure 27. As in Node + 4 mid-split case, we are making significant savings in TCP, and at the same time improving the throughput by 10-15%.



**Figure 27 – Node + 4 high-split power allocation and Rate for channel stacking**

## 6. Conclusion and Future Work

This paper shows that there are significant benefits to be gained from careful allocation of channels to the CMs and optimization of available limited transmit power in a ESD system.

For a given channel allocation, considerable throughput gain, in order of 10%, is achievable with closed-loop power optimization. Furthermore, carefully combining the channel allocation and optimal power distribution can significantly reduce the required TCP for the node and the network amps.

The additional headroom created in transmit power budget can be exercised to improve the range of the network. On the other hand, savings made to overall node power consumption budget can enable more flexible fiber deep deployment options, such as soft FMA and edge compute.

The above benefits are applicable to network topologies across the board, be it passive Node + 0, fiber deep, or conventional Node + X.

Further work is needed to quantify the effect of relaxing some of the network design principles, such as unity gain.

## Abbreviations

|       |                                               |
|-------|-----------------------------------------------|
| C/CSO | carrier to composite second order distortion  |
| C/CTB | carrier to composite triple beat              |
| CM    | cable modem                                   |
| DAA   | distributed access architecture               |
| DS    | downstream                                    |
| ESD   | extended spectrum DOCSIS                      |
| FDD   | frequency division duplexing                  |
| FMA   | flexible MAC architecture                     |
| HFC   | hybrid fiber-coax                             |
| MER   | modulation error ratio                        |
| MTPR  | missing tone power ratio                      |
| NF    | noise figure                                  |
| OFDM  | orthogonal frequency division multiplexing    |
| PA    | power amplifier or power amplification        |
| PMA   | profile management application                |
| PNM   | proactive network maintenance                 |
| PON   | passive optical network                       |
| PSD   | power spectral density                        |
| RMD   | remote MAC device                             |
| RPD   | remote PHY device                             |
| ISBE  | International Society of Broadband Experts    |
| SCTE  | Society of Cable Telecommunications Engineers |
| TCP   | total composite power                         |
| US    | upstream                                      |
| vCMTS | virtual Cable Modem Termination System        |

# Bibliography & References

- [1] IEEE, “IEEE P802.3ca/D1.4 Draft Standard for Ethernet Amendment: Physical Layer Specifications and Management Parameters for 25 Gb/s and 50 Gb/s Passive Optical Networks,” 2018, IEEE 802.3ca/D1.4, 27 November 2018.
- [2] Wang, Eric, “Draft text for G.mgfast-PHY,” 209, ITU Draft Recommendation Q4/15-TD48 (190121).
- [3] CM-SP-PHYv3.1, “Data-Over-Cable Service Interface Specification DOCSIS 3.1, Physical Layer Specification,” CableLabs, 2013
- [4] CM-SP-PHYv4.0, “Data-Over-Cable Service Interface Specification DOCSIS 4.0, Physical Layer Specification,” CableLabs, 2020.
- [5] Ed Harstead, Doutje van Veen, Vincent Houtsma, and Pascal Dom, “Technology Roadmap for Time-Division Multiplexed Passive Optical Networks (TDM PONs),” *Journal of Lightwave Technology*, vol. 37, no. 2, pp. 657–664, 2019.
- [6] Ed Harstead and Randy Sharpe, “Forecasting of access network bandwidth demands for aggregated subscribers using monte carlo methods,” *IEEE Communications Magazine*, vol. 53, no. 3, pp. 199–207, 2015.
- [7] Jakob Nielsen, “Nielsen’s Law of Internet Bandwidth,” 2018.
- [8] Cloonan, Tom and Al-Banna Ayham and O’Keeffe Frank, “Using DOCSIS to Meet the Larger BW Demand of the 2020 Decade and Beyond,” 2016, Arris White Paper.
- [9] Partha P Mitra and Jason B Stark, “Nonlinear limits to the information capacity of optical fibre communications,” *Nature*, vol. 411, no. 6841, pp. 1027, 2001.
- [10] NetCommWireless, “CTTdp Unit (4 ports),” July 2016, Specification Sheet NDD-4200 R1.
- [11] Krapp, Steven, “Virtual Fiber - 100 Gbps over Coax,” in *SCTE 2017 Fall Technical Forum*. IEEE, 2017.
- [12] J.M. Cioffi, “A multicarrier primer,” *ANSI T1E1*, vol. 4, pp. 91–157, 1991.
- [13] Thomas M Cover and Joy A Thomas, *Elements of information theory*, John Wiley & Sons, 2012.
- [14] Rainer Strobel, *Channel Modeling and Physical Layer Optimization in Copper Line Networks*, Springer, 2019.
- [15] Danny Van Bruyssel, Yannick Lefevre, Vladimir Oksman, Rainer Strobel, Miguel Peeters, and Dong Wei, “G.mgfast: Working text for advanced coding scheme,” January 2019, Contribution ITU-T Q4/16-C17-R1 (190121).
- [16] Rob Howald, “Evaluation of In-Home Architecture variants for FDX and Rationale,” April 2017, Comcast FDX f2f Meeting.
- [17] Rainer Strobel and Thushara Hewavithana, “Power Spectrum Optimization for Capacity of the Extended Spectrum Hybrid Fiber Coax Network”, *IEEE ICASSP*, 2020.

# **The Headend Evolution: Design Considerations for Deploying vCCAP and Other VNFs**

A Technical Paper prepared for SCTE•ISBE by

**Patricio Sebastian Latini**  
Regional VP - CALA  
CASA Systems  
100 Old River Rd. – Andover, MA  
+1 (305) 504-9250  
[patricio.latini@casa-systems.com](mailto:patricio.latini@casa-systems.com)

# Table of Contents

| Title                                                    | Page Number |
|----------------------------------------------------------|-------------|
| 1. Introduction .....                                    | 4           |
| 2. Virtualization .....                                  | 4           |
| 2.1. Virtualization Introduction.....                    | 4           |
| 2.2. The Hypervisor .....                                | 5           |
| 2.3. Virtual Machines .....                              | 6           |
| 2.4. Network Function Virtualization.....                | 7           |
| 2.5. Management and Orchestration.....                   | 7           |
| 3. Containers .....                                      | 8           |
| 3.1 – Containers and Microservices.....                  | 10          |
| 3.2 – Containers Orchestration.....                      | 11          |
| 3.3 – Continuous Integration/Continuous Deployment ..... | 12          |
| 4. The Clouds .....                                      | 13          |
| 4.1. Private Cloud.....                                  | 13          |
| 4.2. Public Cloud .....                                  | 14          |
| 4.3. Hybrid Cloud .....                                  | 15          |
| 4.4. Platform as a Services (PaaS).....                  | 15          |
| 5. Virtualizing the CCAP.....                            | 17          |
| 5.1. Control and User Plane Separation (CUPS).....       | 17          |
| 5.2. Virtual CCAP Core and Remote PHY.....               | 18          |
| 5.3. Multi-access Edge Computing (MEC) .....             | 20          |
| 5.4. Virtual MAC Manager and Remote MAC/PHY.....         | 21          |
| 6. Future and Convergent Core Functions.....             | 22          |
| 7. Conclusion .....                                      | 22          |
| Abbreviations.....                                       | 23          |
| Bibliography & References .....                          | 24          |

## List of Figures

| Title                                                       | Page Number |
|-------------------------------------------------------------|-------------|
| Figure 1 – Virtual Machines .....                           | 5           |
| Figure 2 - Hypervisor Types.....                            | 6           |
| Figure 3 - Network Function Virtualization examples .....   | 7           |
| Figure 4 - NFV Paradigm.....                                | 8           |
| Figure 5 - NFV MANO Framework .....                         | 8           |
| Figure 6 - Containers Architecture .....                    | 9           |
| Figure 7 - Monolithic vs Microservices Architectures .....  | 11          |
| Figure 8 - Kubernetes Architecture .....                    | 12          |
| Figure 9 - CI/CD Process.....                               | 13          |
| Figure 10 - Cloud Services Models .....                     | 15          |
| Figure 11 - Cloud Native Hybrid Architecture .....          | 16          |
| Figure 12 - Integrated CCAP vs virtual CCAP Functions ..... | 18          |
| Figure 13 - vCCAP Node to Container Mapping.....            | 19          |
| Figure 14 - vCCAP Kubernetes Managed Architecture.....      | 19          |
| Figure 15 - vCCAP Container Based Redundancy .....          | 20          |
| Figure 16 - vCCAP Deployments Models .....                  | 20          |

|                                                            |    |
|------------------------------------------------------------|----|
| Figure 17 - Integrated Remote MAC-PHY Device .....         | 21 |
| Figure 18 - Decoupled Remote MAC - Remote PHY Devices..... | 21 |
| Figure 19 - Fixed Mobile Convergence Evolution .....       | 22 |

## List of Tables

| <b>Title</b>                            | <b>Page Number</b> |
|-----------------------------------------|--------------------|
| Table 1 - Hypervisor Examples .....     | 6                  |
| Table 2 - Container Engines .....       | 9                  |
| Table 3 - Private Cloud Platforms ..... | 13                 |
| Table 4 - Cloud Service Models .....    | 14                 |
| Table 5 - PaaS Platforms.....           | 16                 |

# 1. Introduction

The telecommunications industry is well underway to moving to virtualizing network functions, and the cable industry is no exception. This paper focuses on providing understanding of virtualized solutions and technology, by analyzing design aspects, capacity planning and architecture evolutions of a CCAP virtual function.

Initially, different virtualization technologies such as virtual machines vs containers are compared, together with continuous integration and deployment as a key element for the required agility to deploy new services and network functions. It is important to mention how the evolution to a distributed computing model and particularly how Mobile Edge Computing (MEC) and network slicing will shape future networks.

Next, a network architecture design is presented focusing on the evolution to separate control and user planes and the impact on network traffic, how they align with distributed access architectures in two flavors, Remote PHY and Remote MAC PHY, and how the vCCAP function could split in two new logical functions in the near future.

Lastly, a set of conclusions is presented to help cable operators better understand the requirements, design options and tradeoffs of vCCAP and DAA deployments in general for the next two to three years and how this decision will impact on the deployment of other related VNFs such as BNG (Broadband Network Gateway) or an EPC (Evolved Packet Core).


## 2. Virtualization

### 2.1. Virtualization Introduction

Generally speaking, Virtualization is defined as running one or multiple instances of a computer system on a layer which is abstracted from the hardware. Each abstracted instance is called a virtual instance. Over the year's virtualization evolved from being just a way of running more than one operating system on a desktop computer at the same time, at the expense of noticeable performance degradation to now being a ubiquitous technology in the server world.

Virtualization offered the ability to run different operating systems, and multiple instances of each. This concept allowed a large system to be split into multiple smaller ones, and hence where now a server could be used to run multiple applications or services while allowing each of them to run completely isolated of the other like it was running in its own dedicated server (Morabito, Cozzolino, Ding, Bejar, & Ott, 2018) - Figure 1.






**Figure 1 – Virtual Machines**

## **2.2. The Hypervisor**

As seen in Figure 1, in order to run multiple virtual machines in physical hardware, a software component is required. This software component is called a hypervisor, a program which takes care of allocating the available resources to each of the virtual machines. Any program run under the hypervisor should exhibit an effect identical with that demonstrated if the program had been run on the original machine directly (Popek & Goldberg, 1974). There are two types of hypervisors as seen in Figure 2. Type one hypervisors are operating systems themselves and run the guest virtual machines directly. Type two hypervisors need to run on a pre-installed operating system and run as application that can be stopped or started.



**Figure 2 - Hypervisor Types**

Some examples of commercial and open source hypervisors are listed in Table 1.

**Table 1 - Hypervisor Examples**

| Type 1                       | Type 2                                        |
|------------------------------|-----------------------------------------------|
| VMware ESX and ESXi          | VMware Workstation                            |
| Microsoft Hyper-V            | Oracle VM VirtualBox                          |
| Citrix XenServer (Xen Based) | Red Hat Enterprise Virtualization (kvm Based) |
| Oracle VM (Xen Based)        |                                               |

### 2.3. Virtual Machines

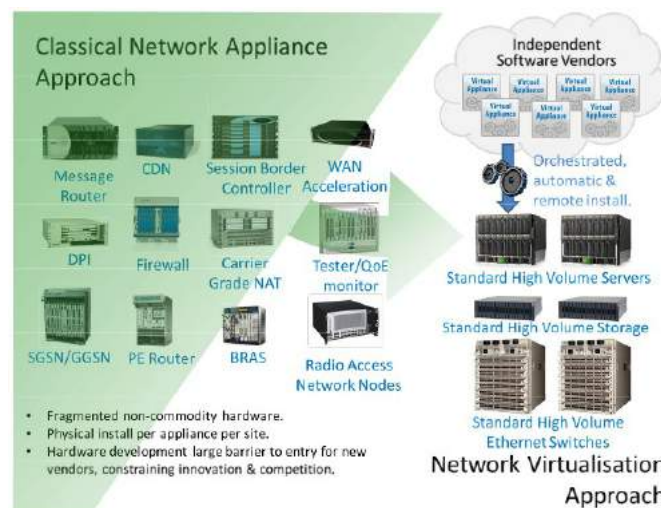
A virtual machine is a term that dates back to 1974 and is defined as an efficient and isolated duplicate of a real machine, where a piece of software provides an environment which is essentially identical to the original machine and programs can run with only minor decreases in speed (Popek & Goldberg, 1974). Virtual machines or guests can have access to the resources available on the host system. The host system can provide computing power, memory, disk space and access to the network interface cards. In the early days of x86 virtualization, hypervisors needed to fully emulate the behavior of the virtual machine including the CPU instruction set, causing a big overhead of computing power and significantly affecting performance of the whole

virtualized system. In 2006, Intel and AMD introduced hardware implemented virtualization extensions (VT and SVM) which allowed the hypervisors to directly access the CPU instructions thus making x86 virtualization possible in terms of performance (Adams & Agesen, 2006). Depending on the expected quality of service, sometimes those resources could or could not be oversubscribed.

## 2.4. Network Function Virtualization

In 2013 ETSI coined the term Network Function Virtualization (NFV). At that time telecommunications operators' networks were populated with a large and increasing variety of proprietary hardware appliances. Launching a new service required even more hardware and the space and power requirements to integrate those boxes was becoming a real challenge. At the same time energy costs were increasing and also hardware obsolescence cycles were becoming shorter, increasing CAPEX costs (European Telecommunications Standards Institute, 2012). At the same time off-the-shelf server computing power was increasing with the costs of those servers being reduced.

Network Functions Virtualization's goal is to address these problems by providing standardization to the virtualization technology in order to consolidate different network equipment types into industry standard off the shelves servers, switches and storage (European Telecommunications Standards Institute, 2012) as seen in Figure 3.

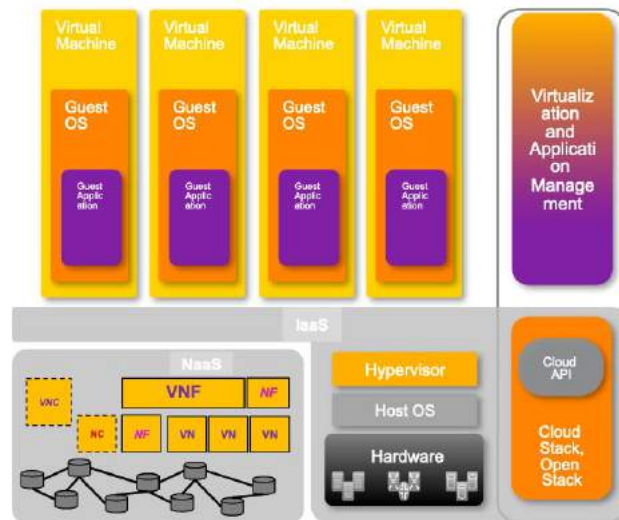


**Figure 3 - Network Function Virtualization examples**

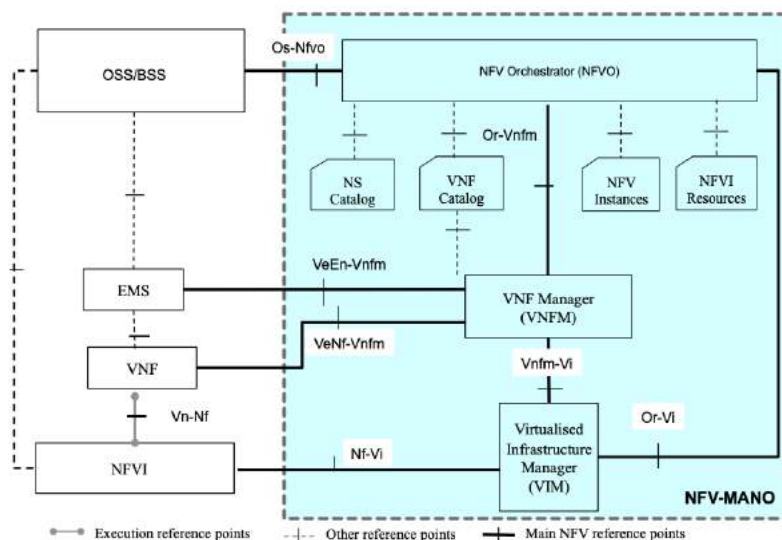
## 2.5. Management and Orchestration

Network function virtualization (NFV) changes how networks are managed. Hundreds of network functions running on a server farm can easily make operations very complex. For that reason the European Telecommunications Standards Institute (ETSI) started in 2013, an effort to define the framework shown in Figure 4 for NFV management from the initial set-up, to day-to-day operations, which is called NFV MANO (Management and Network Orchestration). This is the framework for the management and orchestration of all resources in a data center for virtualized functions, those resources include: compute, networking, storage, and virtual

machines (VM) as seen in Figure 5. The main goal of MANO is to allow flexible on-boarding. (European Telecommunications Standards Institute, 2012)



### Figure 4 - NFV Paradigm



### Figure 5 - NFV MANO Framework

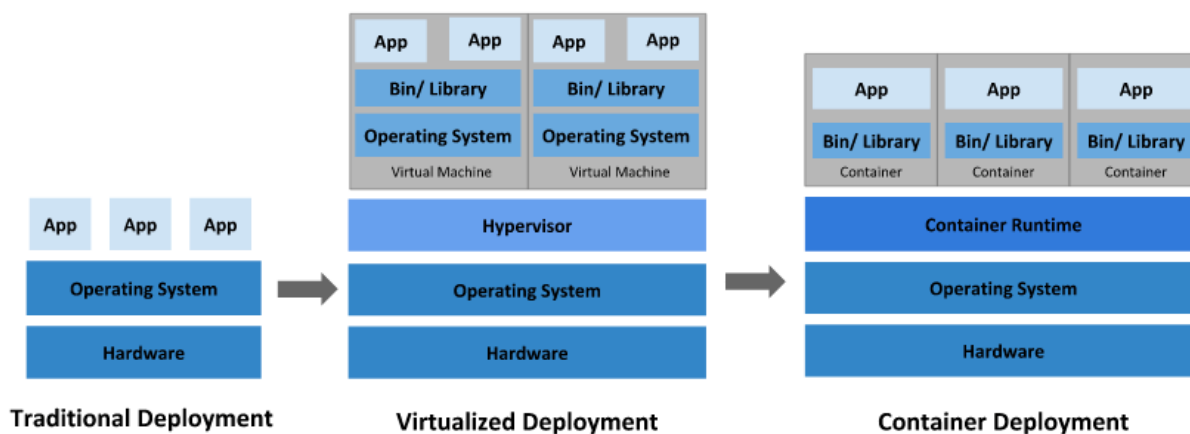
### 3. Containers

As VMs were deployed, organizations started to run applications on separate virtual machines in order to dedicate their own resources (CPU, memory, storage). However, each application had to have a separate operating system running in each virtual machine.

At scale virtual machines may take up a lot of system resources, and a base operating system installation may take several gigabytes of storage. For example, running a single webserver on a virtual machine may require several gigabytes of operating system and libraries, while the application itself only takes a few megabytes of storage and memory. Multiply that by a big number of webserver and other applications, and the overhead of virtualization becomes very noticeable.

At the same time, with the wide adoption of Linux as an operating system, and a big share of network applications using it as its underlying platform, an interesting idea was born: Why can't applications share the same operating system and be virtualized at the operating system level instead of at the hardware level?

A container is an isolated environment from the operating system host that runs an application by virtualizing it. This allows it to create multiple application workloads on a single OS instance. The kernel of the host operating system provides the required components for running the different functions of an application, separated into containers as shown in Figure 6. Containers run isolated tasks from each other, and an application cannot harm the host machine nor come in conflict with other apps running in other containers. (Simic, 2019)



Retrieved from: <https://kubernetes.io/docs/concepts/overview/what-is-kubernetes/>

**Figure 6 - Containers Architecture**

Docker is the most common container engine now in the market, however there are several other options for container engines shown in Table 2.

**Table 2 - Container Engines**

|                           |
|---------------------------|
| Docker                    |
| Mesos                     |
| LXC                       |
| OpenVZ                    |
| Java Container            |
| Windows Server Containers |

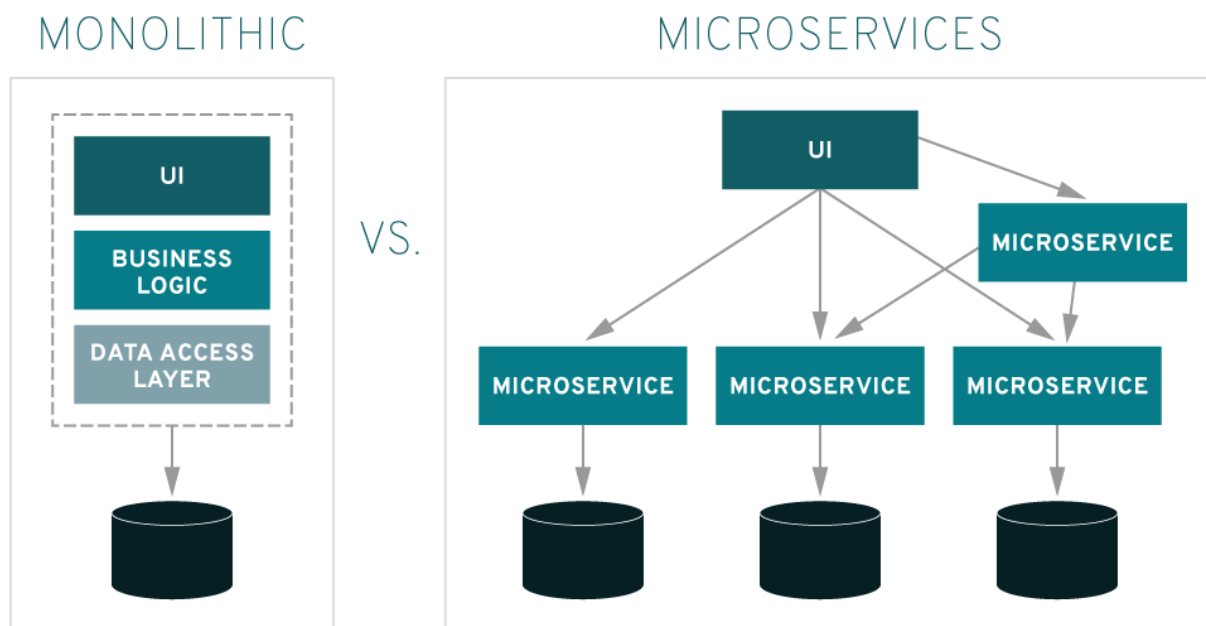
### 3.1 – Containers and Microservices

Applications are easy to develop using a monolithic approach, but as the size of the application and its user base grows it becomes very complex to scale up. Monolithic applications can have up to hundreds of different services tightly coupled, which makes it very complex for different development teams to handle their coordination. This is the reason why most of the software industry is moving to a new paradigm, known as microservice architecture (Yu, Silveira, & Sundaram, 2016).

In a microservice approach, an application consists of several services, each independent of the other. Each service does a specific function which is developed and deployed independently from the others. Services transfer data or information to other services using a standardized communication protocol as seen in Figure 7.

Some of the benefits of microservices are easier automated testing, flexible deployment models and increased overall resiliency; however all this comes at the expense of careful planning of the architecture and increased R&D investment, mainly given to the following factors:

- 1) Since everything is an independent service, careful handling of requests traveling between the modules is required. This generates the need of designing in advance the APIs that the microservices will use to communicate between them its maintenance across version changes.
- 2) Testing a microservices based application can be complex. In a monolithic approach, the application just needs to be launched and validate its connectivity with the underlying services such as database or webserver. With microservices, each service needs to be confirmed to be working properly in advance before integration testing of all the microservices as an application can occur.



## Figure 7 - Monolithic vs Microservices Architectures

As has been shown before, a container is just a way of deploying and running a program or process isolated from others sharing a common operating system. One could have one big monolithic application running as a container and in the other side there could be a big number of microservices running on a bare operating system. However, in the real world these independent approaches have been very complementary, as containers are the portable code envelope and big applications could be decomposed into many microservices that can be independently deployed.

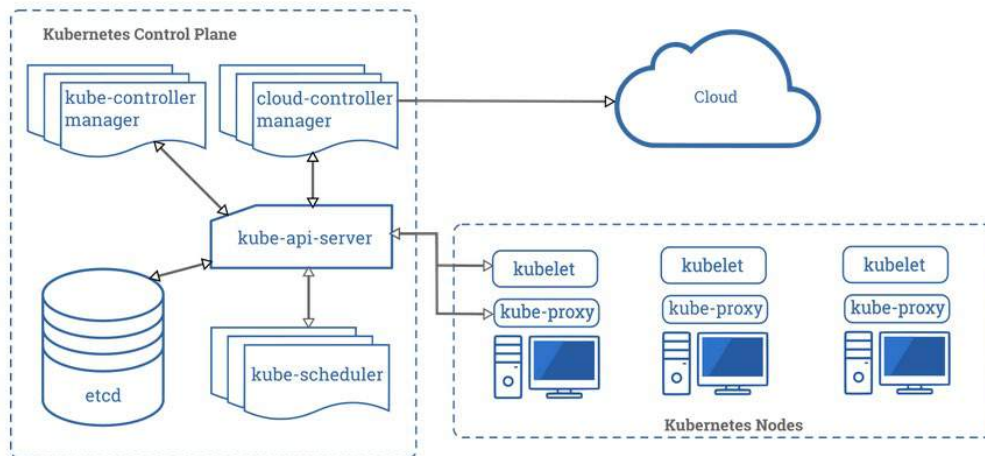
DevOps processes are the pillar of modern applications, and support running multiple parts of applications independently in different microservices, with much greater control over their life cycles.

### 3.2 – Containers Orchestration

Container orchestration is the process of automating the deployment, management and scaling of containers. As mentioned before, service providers will need to deploy hundreds or thousands of containers and for that reason an automation process for container orchestration is required. Container orchestration can help to deploy the same application across multiple environments and microservices in those containers and make it easier to orchestrate different types of services such as storage, networking, and security. (Redhat, n.d.)

There are several orchestration tools for managing containers at scale together with its lifecycle management. The most popular is called Kubernetes, however Docker Swarm, and Apache Mesos are also well-known orchestration tools.

Kubernetes is an open source orchestration tool that was developed by Google. Kubernetes orchestration allows the deployment of applications that are supported by multiple containers, running those containers in clusters of servers and verifying their health and scale-in or scale-out as required by capacity demands. In general, Kubernetes eliminates most of the manual processes associated with deploying complex containerized applications. The Kubernetes architecture is presented in Figure 8.



Retrieved from: <https://kubernetes.io/docs/concepts/overview/what-is-kubernetes/>

**Figure 8 - Kubernetes Architecture**

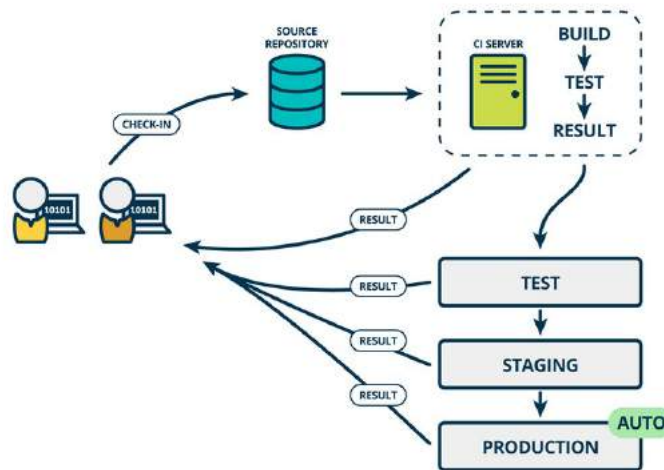
### 3.3 – Continuous Integration/Continuous Deployment

Continuous software engineering is an area which refers to the development, deployment and feedback from software and customers in a very rapid cycle (Fitzgerald & Stol, 2017). The continuous software engineering process can be split into two main areas

Continuous Integration (CI) is a widely established development practice in the software industry where teams integrate and merge development code very frequently, for example multiple times per day. CI allows software developers to have shorter release cycles and improve software quality (Fitzgerald & Stol, 2017). Many of the processes involved in CI are automated.

Continuous Deployment (CD) is a practice that goes a step further and automatically and continuously deploys the application to a production environment as seen in Figure 9. That environment can be within the same company or with an external customer. In the latter case, Continuous Delivery (CDE) may apply, where the application is only delivered but not automatically turned into production and requires some manual intervention before being deployed.





Retrieved from: <https://www.mindtheproduct.com/what-the-hell-are-ci-cd-and-devops-a-cheatsheet-for-the-rest-of-us/>

**Figure 9 - CI/CD Process**

Two common software tools for this process are GIT as the source code repository, and Jenkins as the CI Server. Jenkins builds the application from the code repository, automatically tests, and delivers the containers to a container repository.

## 4. The Clouds

### 4.1. Private Cloud

According to the National Institute of Standards and Technology (NIST), a private cloud is defined as infrastructure which is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises (National Institute of Standards and Technology , 2011).

The main advantage of a private cloud is that resources are not shared. A private cloud is best for businesses with dynamic needs that require direct control over their computing resources, in general to meet security or regulatory requirements.

Private clouds also have a few disadvantages as increased automation and user self-service can bring added complexity. These technologies require teams to rearchitect data centers and use extra management tools and can result in increased staff and capital expenditures to acquire the infrastructure to support it.

When a private cloud is properly architected and implemented, the organization can benefit from self-service and scalability, and the ability to launch or optimize computing resources on demand. In Table 3 there is a list of the main private cloud platforms in the market. It is important to mention that Openstack-based platforms are provided by several mainstream vendors.

**Table 3 - Private Cloud Platforms**

|           |
|-----------|
| Openstack |
|-----------|

## 4.2. Public Cloud

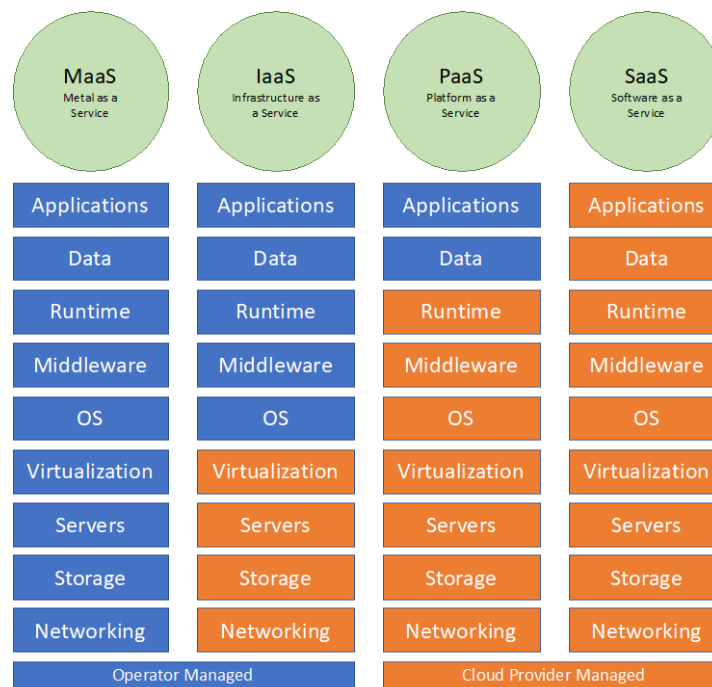
According to the National Institute of Standards and Technology (NIST), a public cloud is the cloud infrastructure that is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. (National Institute of Standards and Technology , 2011)

From a service provider standpoint, a public cloud is a platform in which a third-party service provider makes available computing resources which can be software applications, virtual machines (VMs) and complete enterprise-grade infrastructures over the public Internet. This public cloud service provider owns the data centers where customers' services run. Service providers take care of all the infrastructure maintenance and provide connectivity access to applications and data. (IBM, 2020)

There are four main models of cloud services as shown in Table 4 together with examples of cloud providers for each. Each cloud service model has a different management complexity compared to the other, as shown in Figure 10.

**Table 4 - Cloud Service Models**

|                                    |                                                                                       |
|------------------------------------|---------------------------------------------------------------------------------------|
| Software as a Service (SaaS)       | Google Apps, Dropbox, Salesforce, Cisco, Concur, GoToMeeting, Slack                   |
| Platform as a Service (PaaS)       | Google App Engine, Force.com, Redhat Openshift                                        |
| Infrastructure as a Service (IaaS) | Amazon Web Services (AWS), Microsoft Azure, Google Compute Engine, Rackspace          |
| Metal as a Service (MaaS)          | Amazon Web Services Bare Metal, Microsoft Azure Bare Metal Servers, Google Bare Metal |



**Figure 10 - Cloud Services Models**

### 4.3. Hybrid Cloud

According to the National Institute of Standards and Technology (NIST), a hybrid cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds) (National Institute of Standards and Technology , 2011).

### 4.4. Platform as a Services (PaaS)

There are several common core microservices which could be shared in either a private or public cloud, without having to replicate them in each application. It may be desirable that service providers have them as a shared PaaS services. Some of them are listed below.

- **Message Queue.** Responsible for delivering messages from one service instance to another. The message queue will support different messaging patterns including request/response, broadcast, pub/sub model.
- **Session Database.** Responsible for storing session state in a common place, which enables stateless processing.
- **Service Discovery.** Responsible for service registration, service health monitoring and service state notification. Note some platforms like Kubernetes provide built-in service discovery.
- **Configuration store.** Responsible for maintaining persistent configuration for the containers. The service configuration will be modeled as a tree structure. Each micro service will have its own subtree

- **Logging.** Responsible to collect logs from all micro-service instances. This service provides a single access point to view the logs.
- **Analytic and Visualization.** Responsible for providing in-depth knowledge of the platform status and presents the result in an easy-to-understand graphical form.
- **Management API.** Responsible for providing RESTAPI/CLI to configure the platform and provide service statistics or session information.

An example implementation of such architecture is shown in Figure 11 with a list of some of the most common software packages for them listed in Table 5.

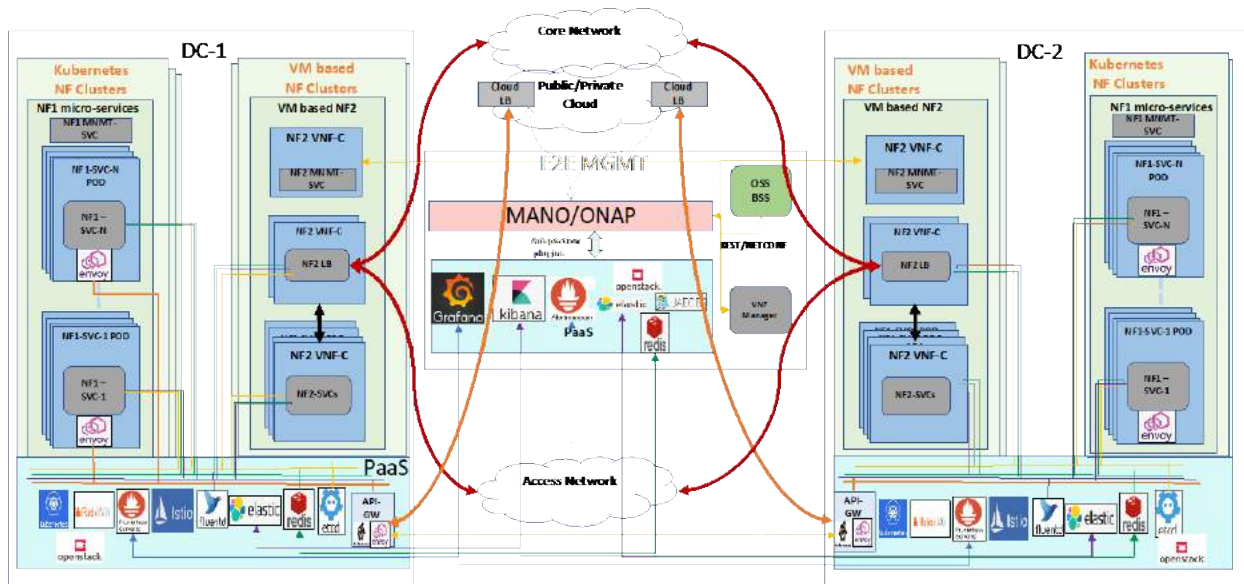













Figure 11 - Cloud Native Hybrid Architecture

Table 5 - PaaS Platforms

|                                                                                     |                 |                                                |
|-------------------------------------------------------------------------------------|-----------------|------------------------------------------------|
|  | Openstack       | VM Orchestration                               |
|  | Kubernetes      | micro-service container orchestration          |
|  | Docker Registry | NF micro-service Container image               |
|  | Helm            | Application packaging, deployment and upgrades |
|  | RabbitMQ        | Intra-NF micro-services message bus            |

|                                                                                   |                          |                                                 |
|-----------------------------------------------------------------------------------|--------------------------|-------------------------------------------------|
|  | ETCD                     | NF micro-service specific configuration store   |
|  | Istio/Envoy              | Service Mesh                                    |
|  | Fluentd/ELK Stack/Jaeger | Logging/tracing                                 |
|  | Prometheus/Grafana       | Monitoring and Alerting                         |
|  | Redis-DB                 | NF micro-service specific state and stats store |
|  | Calico/Multus            | Container Networking                            |

## 5. Virtualizing the CCAP

### 5.1. Control and User Plane Separation (CUPS)

Cable Operators have historically used CMTS as Edge equipment which has several functions like per-subscriber session, policy enforcement and data forwarding in the same box. With the advent of Control Plane User Plane Separation (CUPS), CMTS functions could be decomposed and disaggregated such that the User Plane Function (UPF) could be deployed in a distributed manner, and the Control Plane Function (CPF) could be deployed in a centralized manner depending on the level of scale and aggregation needed. (Asati & Bernstein, 2019)

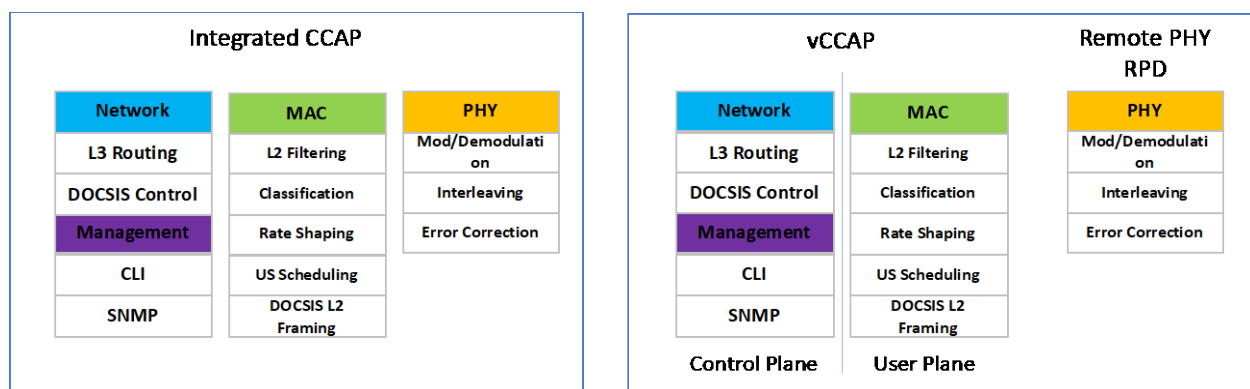
The deployment of new services, such as 4K video, IoT, etc., and increasing numbers of home broadband service users present some new challenges for broadband routers such as:

- Low resource utilization: The traditional CCAP acts as both a gateway for user access authentication and an IP network's Layer 3 edge. The nature of the tightly coupled control plane and forwarding plane makes it difficult to achieve the optimum performance of either of the planes.
- Complex management and maintenance: Due to the large numbers of traditional CCAP instances a network must have, each device must be configured one at a time when deploying global service policies. As the network expands and new services are introduced, this deployment mode will cease to be feasible as it is unable to manage services effectively and rectify faults rapidly. (Hu, et al., 2018)

To address these challenges, a cloud based BNG with CU separation conception is defined in (Broadband Forum, 2018) however the same idea is applicable to a CCAP box. The main idea of Control-Plane and User-Plane separation is to extract and centralize the user management functions of multiple CCAP devices, forming a unified and centralized control plane (CP). The traditional router's Control Plane and Forwarding Plane are both preserved on CCAP devices in the form of a user plane (UP). Note that the CU separation concept has also be introduced in the 3GPP 5G architecture. (3GPP, 2018)

## 5.2. Virtual CCAP Core and Remote PHY

As mentioned in the previous section, the paradigm of separating control and user planes brings significant benefits to the table. Traditional integrated CCAP boxes running on dedicated hardware were not able to benefit from the approach where control plane, user plane and physical modulation were part of the same hardware box. However, with the adoption of Distributed Access Architectures the paradigm has shifted so that the modulation function has moved to the RPD and now the CCAP core may be an IP-in/IP-out only box. This shift brought the possibility of moving CCAP software functions from dedicated hardware onto commercial off-the-shelves (COTS) servers and generated the term virtual CCAP as seen in Figure 12




**Figure 12 - Integrated CCAP vs virtual CCAP Functions**

Two conclusions can also be drawn from Figure 12:

- The CCAP's different software components can be grouped to exploit the benefits of microservices which were analyzed in section 2
- At the same time, it also makes sense to separate the control and user plane functions, as there is an advantage of potentially being able to put them in separate locations.

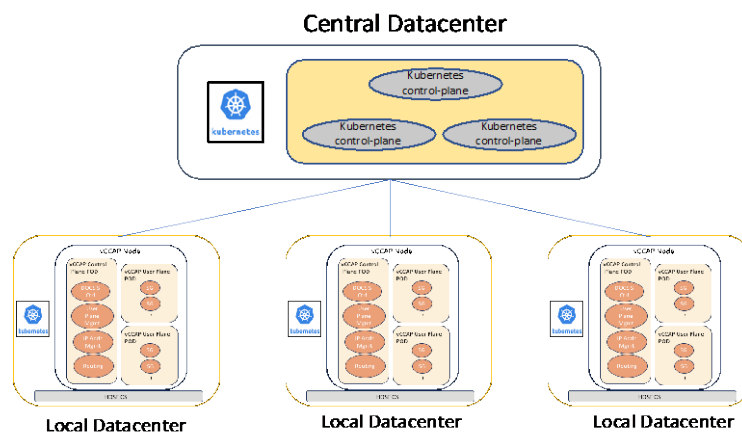
The above points support using containers for the vCCAP user plane and control plane functions. Now one could think that a user plane container can be the equivalent of integrated CCAP MAC card which typically serves 6 to 8 HFC serving groups, but in reality, that constraint does not exist anymore. A user plane container can serve a single serving group as seen in Figure 13; and at the same time scale to manage multiple user plane functions.



**Figure 13 - vCCAP Node to Container Mapping**

No matter how many serving groups per container one decides to use, it is evident that the quantity of containers will be significant, so as analyzed in section 2, a container orchestration engine would be also mandatory in order to manage the growing number of containers. However the benefits of container orchestration are not only related to the onboarding of the containers but also how they can handle dynamic scale-in and scale-out of containers based on server utilization or network traffic, or even on redundancy management in the case of software or hardware failures.


As an example, architecture, a container farm is shown in Figure 14 using Kubernetes as the container orchestrator. Kubernetes relies on control-plane and worker nodes; the control-planes take care of the onboarding, deployment and monitoring of the groups containers or PODs which are dynamically deployed in the worker nodes on multiple datacenters on the operator's network.



**Figure 14 - vCCAP Kubernetes Managed Architecture**

As mentioned before, Kubernetes can monitor health status of the running pods. In the example in Figure 15, if it detects a failure on the hardware or software of a user plane pod, it can switch its operation to a redundant POD.

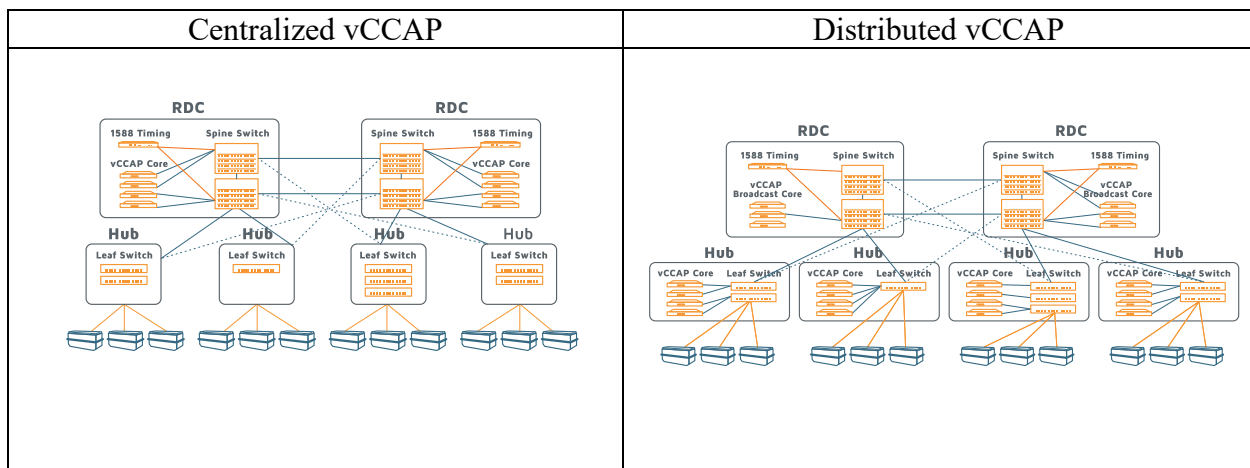




**Figure 15 - vCCAP Container Based Redundancy**

Perhaps one of the most important advantages of deploying virtualized CCAP is the flexibility to support different types of architectures, going from a totally distributed model where the computing power can be in the hubsites (which can be repurposed as data centers) to a totally centralized model where the computing power is a main datacenter, to any combination in-between. Figure 16.

Each model has its advantages and disadvantages depending on space availability, bandwidth consumption and server usage efficiency, however having the ability to mix them and use the most efficient model for each case is one of the key highlights of vCCAP. This is known as a Hybrid model.



**Figure 16 - vCCAP Deployments Models**

### 5.3. Multi-access Edge Computing (MEC)

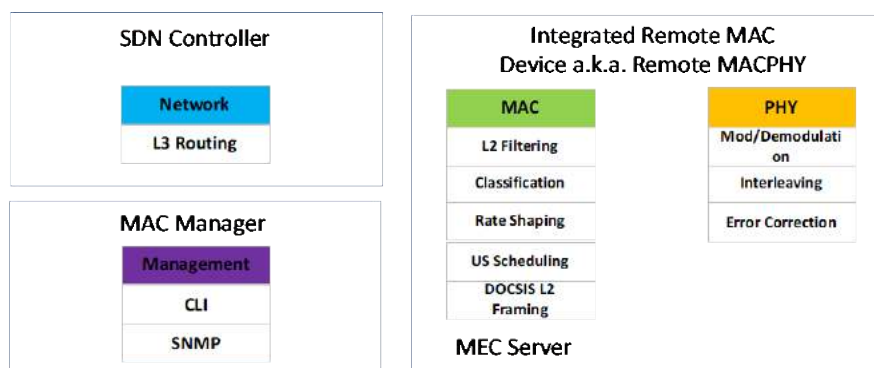
Multi-access edge computing (MEC) is a cloud environment located at the edge of the network, in close proximity to end users and coupled with the service provider's network infrastructure. Even before 5G is rolled out, current fixed and mobile networks can already enable support for these challenging use cases by using MEC technology. MEC is able to offer low latency and high bandwidth, and, in addition allows services to be deployed in different industrial premises such as road infrastructure, airports, and factories, bringing computing power where it is needed




most. MEC is a key enabler for low latency IoT technology and having this computing power on a node in the network can provide a big number of extra benefits apart of the network function virtualization (Porambage, Okwuibe, Liyanage, Taleb, & Ylianttila, 2018). ETSI Industry Specification Group MEC were the pioneers in creating a standardized computing platform for mobile networks applying network edge related use cases. (Giust, et al., 2018)

#### 5.4. Virtual MAC Manager and Remote MAC/PHY

The Remote MAC-PHY technology moves both the DOCSIS MAC and PHY layers down to the Remote/Fiber Node. The link between the Headend and the node is essentially a Layer 2 connection using Ethernet (Cable Television Laboratories, Inc, 2015). This new device is called the Remote MAC Device (RMD), where this device can be integrated with a Remote PHY Device (RPD) as in Figure 17 or both functions can be in different devices as in Figure 18.



**Figure 17 - Integrated Remote MAC-PHY Device**



**Figure 18 - Decoupled Remote MAC - Remote PHY Devices**

It is particularly relevant to highlight that as mentioned in section 5.2, the MAC function is the user plane of a virtual CCAP, which could be run as a container in a cloud native platform. So, in this architecture the Remote MAC Device doesn't need to be more than a Multi-access Edge

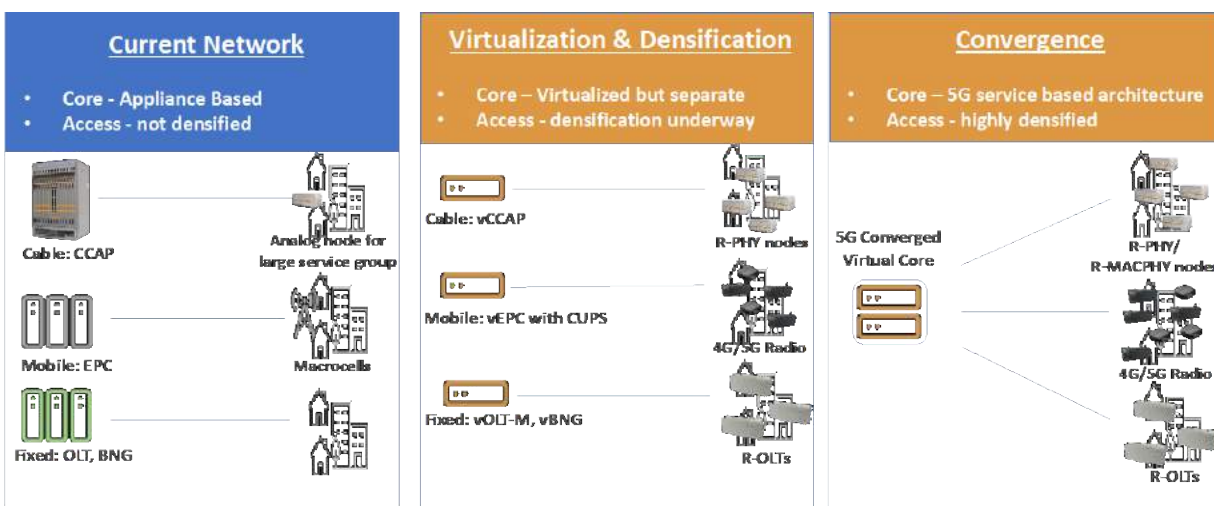
Computing server running on a hardened node enclosure, which is automatically managed by container orchestration engine such as Kubernetes.

At this stage the only remnant of the CMTS is the MAC manager which is a centralized software function which provides management access such as command line interface and monitoring data of the CCAP network function running on the remote devices.

## 6. Future and Convergent Core Functions

Cable Operators can use this technology shift to embrace Fixed Mobile Convergence (FMC) by putting edge network functions at the hub sites which will become virtualization datacenters. As mentioned before, CUPS could enable a common user plane function for cable access, FTTH access and mobile access distributed, while the control plane could be located in centralized datacenters. This is aligned with the decomposed vCCAP and RPHY trend that the industry is following (Asati & Bernstein, 2019). Virtual BNGs and 5G Virtual EPCs can run in the shared cloud environment providing a high level of optimization and integration.

In Figure 19, a potential evolution path for core access networks is presented, where the virtualization of the core and densification of the access will be the first step towards a real Fixed and mobile converged network.



**Figure 19 - Fixed Mobile Convergence Evolution**

## 7. Conclusion

This paper described the state of current virtualization and containerization technologies, discussing benefits and disadvantages applied to cloud technologies. A review of virtual and private cloud was done with a focus on the implementation of virtual network functions and how a continuous integration and deployment approach will help in making networks more agile and flexible.

Next, an architecture proposal was presented for deploying a virtual CCAP network function leveraging the benefits of cloud native platforms together with the benefits of control and user

plane separation. A brief discussion of Multi Access edge computing aligned with the transition path to Flexible MAC Architectures and remote MAC-PHY was shown.

Lastly, a vision for future Fixed Mobile Convergence is presented with the evolution through virtualization of the network functions and densification of the access network.

As a final conclusion: HFC networks need to evolve into adopting the benefits provided by virtualization, containerization and cloud technologies. As reviewed during the paper, those technologies provide not only big improvements in efficient usage of the infrastructure but more importantly bring the ability to deploy new services in a quick and agile manner, something that will be a key factor for cable operators to continue its success in the new world of Fixed Mobile Converged networks.

## Abbreviations

|        |                                                 |
|--------|-------------------------------------------------|
| BNG    | Broadband Network Gateway                       |
| CAPEX  | Capital Expenditures                            |
| CCAP   | Cable Converged Access Platform                 |
| CD     | Continuous Deployment                           |
| CI     | Continuous Integration                          |
| CMTS   | Cable Modem Termination System                  |
| CPF    | Control Plane Function                          |
| CPU    | Central Processing Unit                         |
| CUPS   | Control and User Plane Separation               |
| DevOps | Software Development / IT Operations            |
| EPC    | Evolved Packet Core                             |
| ETSI   | European Telecommunications Standards Institute |
| FTTH   | Fiber to the Home                               |
| IaaS   | Infrastructure as a Service                     |
| IoT    | Internet of Things                              |
| IP     | Internet Protocol                               |
| ISBE   | International Society of Broadband Experts      |
| MaaS   | Metal as a Service                              |
| MAC    | Media Access Control                            |
| MANO   | Management and Network Orchestration            |
| MEC    | Multi-Access Edge Computing                     |
| NFV    | Network Function Virtualization                 |
| OS     | Operating System                                |
| PaaS   | Platform as a Service                           |
| HFC    | hybrid fiber-coax                               |
| PHY    | Physical                                        |
| R&D    | Research and Development                        |
| RMD    | Remote MAC Device                               |
| RPD    | Remote PHY Device                               |
| SaaS   | Software as a Service                           |
| SCTE   | Society of Cable Telecommunications Engineers   |

|       |                                         |
|-------|-----------------------------------------|
| UPF   | User Plane Function                     |
| vCCAP | Virtual Cable Converged Access Platform |
| VNF   | Virtual Network Function                |
| VT    | Virtualization Technology               |

## Bibliography & References

- 3GPP. (2018). *System Architecture for the 5G System 3GPP GPP TS 23.501 15.0.0*. 3GPP.
- Adams, K., & Agesen, O. (2006). A comparison of software and hardware techniques for x86 virtualization. *ACM SIGOPS Operating Systems Review*, 40(5).
- Asati, R., & Bernstein, A. (2019). Cable Edge Compute: Transforming Cable Hubs into Application-Centric Clouds. *SCTE Cable-TEC 2019*. New Orleans, LA.
- Broadband Forum. (2018). *Cloud Central Office Reference Architectural Framework TR-384*. Broadband Forum.
- Cable Television Laboratories, Inc. (2015). *Distributed CCAP Architectures Overview Technical Report*. Cable Television Laboratories, Inc.
- European Telecommunications Standards Institute. (2012). Network Functions Virtualisation: An Introduction, Benefits, Enablers, Challenges & Call for Action. *SDN and OpenFlow World Congress*. Darmstadt, Germany.
- European Telecommunications Standards Institute. (2013). ETSI NFV Management and Orchestration - An Overview. *IETF #88*. Vancouver, CA.
- Fitzgerald, B., & Stol, K. (2017). Continuous software engineering: A roadmap and agenda. *The Journal of Systems and Software*, 123, 176–189.
- Giust, F., Sciancalepore, V., Sabella, D., Filippou, M. C., Mangiante, S., Featherstone, W., & Munaretto, D. (2018). Multi-Access Edge Computing: The Driver Behind the Wheel of 5G-Connected Cars. *IEEE Communications Standards Magazine* (September), 66-73.
- Hu, S., Qin, F., Li, X., Chua, T., Eastlake, D., Wang, Z., & Song, J. (2018, October 22). *Architecture for Control Plane and User Plane Separated BNG*. Retrieved July 31, 2020, from <https://tools.ietf.org/id/draft-cuspd-tgtwg-cu-separation-bng-architecture-02.html>
- IBM. (2020, March 3). *Public Cloud Cloud*. Retrieved July 31, 2020, from <https://www.ibm.com/cloud/learn/public-cloud>
- Morabito, R., Cozzolino, V., Ding, A. Y., Beijar, N., & Ott, J. (2018). Consolidate IoT Edge Computing with Lightweight Virtualization. *IEEE Network*.
- Motjaba, S., Babar, M. A., & Zhu, A. L. (2017). Continuous Integration, Delivery and Deployment: A Systematic Review on Approaches, Tools, Challenges and Practices. *IEEE Access*, 5(2017), 3909-3943.
- National Institute of Standards and Technology . (2011, September). *The NIST Definition of Cloud Computing*. Retrieved July 31, 2020, from <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>
- Popek, G. J., & Goldberg, R. P. (1974). Formal Requirements for Virtualizable Third Generation Architectures. *Communications of the ACM*, 17(7), 412-421.
- Porrambage, P., Okwuibe, J., Liyanage, M., Taleb, T., & Ylianttila, M. (2018). Survey on Multi-Access Edge Computing for Internet of Things Realization. *IEEE Communications Surveys & Tutorials*.

- Redhat. (n.d.). *What is container orchestration?* Retrieved July 31, 2020, from <https://www.redhat.com/en/topics/containers/what-is-container-orchestration>
- Simic, S. (2019, April 19). *Containers vs Virtual Machines (VMs): What's the Difference?* Retrieved from <https://phoenixnap.com/kb/containers-vs-vms>
- The Linux Foundation. (n.d.). *What is Kubernetes?* Retrieved July 30, 2020, from <https://kubernetes.io/docs/concepts/overview/what-is-kubernetes/>
- Yu, Y., Silveira, H., & Sundaram, M. (2016). A microservice based reference architecture model in the context of enterprise architecture. *2016 IEEE Advanced Information Management, Communicates, Electronic and Automation Control Conference*. Xi'an, China.

# **Powering The Future: Next Generation Access Networks**

A Technical Paper prepared for SCTE\*ISBE by

**Todd Musat**

Director, Critical Network Infrastructure  
Shaw Communications Inc.  
2728 Hopewell Place NE Calgary AB T1Y 7J7  
1.403.930.9634  
[todd.musat@sjrb.ca](mailto:todd.musat@sjrb.ca)

**Chuck Carroll**

General Partner  
Saras Energy Consulting  
1006 Monroe Ave River Forest, IL 60305  
1.360.201.4205  
[ccarroll@saraspartners.com](mailto:ccarroll@saraspartners.com)

**Lew Rakowsky**, Contributor, Saras Energy Consulting

**Rene Spee**, Contributor, Saras Energy Consulting

# Table of Contents

| <b>Title</b>                                           | <b>Page Number</b> |
|--------------------------------------------------------|--------------------|
| 1. Introduction .....                                  | 3                  |
| 2. Business as Usual Network .....                     | 4                  |
| 3. Approach to DAA .....                               | 5                  |
| 4. DAA Transition.....                                 | 6                  |
| 5. Modeling Energy Usage .....                         | 8                  |
| 5.1. Initial Energy Usage Hypothesis.....              | 8                  |
| 5.2. Energy Model Assumptions and Structure.....       | 9                  |
| 6. Findings from Modeling .....                        | 12                 |
| 6.1. BAU .....                                         | 12                 |
| 6.2. DAA .....                                         | 13                 |
| 6.3. DAA vs. BAU.....                                  | 15                 |
| 7. Outside Plant Power Network Capacity Analysis ..... | 16                 |
| 8. Conclusion.....                                     | 19                 |
| Abbreviations .....                                    | 20                 |
| Bibliography and References .....                      | 21                 |

## List of Figures

| <b>Title</b>                                                                    | <b>Page Number</b> |
|---------------------------------------------------------------------------------|--------------------|
| Figure - 1 Business as Usual Network Deployment.....                            | 4                  |
| Figure 2 Shaw DAA Implementation.....                                           | 6                  |
| Figure -3 Shaw DAA Implementation.....                                          | 8                  |
| Figure -4 DAA vs. BAU Percentage Energy Change Initial Hypothesis Modeling..... | 9                  |
| Figure -5 BAU Energy Growth.....                                                | 12                 |
| Figure -6 DAA Energy Growth.....                                                | 13                 |
| Figure -7 Percentage Difference DAA vs. BAU.....                                | 15                 |
| Figure -8 PS Distribution of Remaining Output VA.....                           | 18                 |
| Figure -9 Distribution PS Remaining Output VA %'s .....                         | 18                 |

## List of Tables

| <b>Title</b>                               | <b>Page Number</b> |
|--------------------------------------------|--------------------|
| Table -1 Non-HFC Device Typical Power..... | 17                 |

## 1. Introduction

As bandwidth capacity continues to grow for existing hybrid fibre coax (HFC), so does the flexibility to utilize new technologies within HFC networks. Establishing practices which support new technology platforms, while fully employing the large investment in the HFC plant is important. As the coaxial cable is capable for shorter distances of carrying high bit rate traffic, MSOs generally implement a fiber deeper approach, where node splits allow fiber penetration further into the network, and closer to the customer. Often, an N+0 approach is considered, which has no active devices on the "last mile" coax between the optical node and the customer premises. However, cost considerations often require compromises, and in many cases, an N+1 (one amplifier between node and customer) or even N+2 approach is implemented — still a vast improvement over the historical N+4 or higher architectures.

HFC networks traditionally terminate node traffic on cable modem termination systems (CMTSs). Integration of CMTSs with other transport functions in the head-end and hub has driven implementation of the converged cable access platform (CCAP). CCAP combines the functions of the CMTS and the edge quadrature amplitude modulators (QAMs), allowing digital data processing prior to analog conversion for transport on the access network. Although CCAP itself drove more power and space efficient HFC edge network solutions, in a fiber deep architecture, the number of nodes will increase significantly, and with it, the requirement for additional CCAP hardware in the hub. This, in turn requires additional space, power and cooling in the facility.

In the distributed access architecture (DAA), the constraints on the headend are ameliorated since the physical and media access control (PHY and MAC) can be placed further in the access network (e.g. directly with the optical node). There are multiple implementation possibilities for DAA. Currently, the most popular and only fully CableLabs specified approach is Remote PHY, where the CCAP PHY layer is moved to the node, while the MAC layer continues to reside in the hub. The facility improvements in terms of space, power and cooling are moderate, as the CMTSs are not being eliminated.

Other possible implementations of DAA are Remote MAC/PHY and Remote CCAP. Here, both physical and MAC layers are moved to the node. Many functions can be virtualized and this in turn can lead to a drastic reduction of facility space, power and cooling requirements.

In 2019, as an integral part of evolving the network to 10G, Shaw began the process of introducing DAA into its network. As with other MSOs, the key driver of this effort was to move CCAP PHY processing from the hub to the node, replacing RF optics with metro ethernet (metroE) optics. The chosen strategy evolved DAA into the network in a controlled manner, implementing DAA in greenfield nodes and node splits, but not proactively in existing brownfield. As this was a different approach to DAA implementation, the company wished to better understand how a migration to DAA in this manner would affect energy usage across their footprint.

In conjunction with Shaw, Saras Energy Consulting worked to assess the energy use impact of the company's approach to DAA across its footprint. The workplan focused on the following investigative objectives:

- Establish a hypothesis about the impact of DAA on facility and outside plant (OSP) energy environments, based on the DAA design and implementation plan as per Shaw standard.

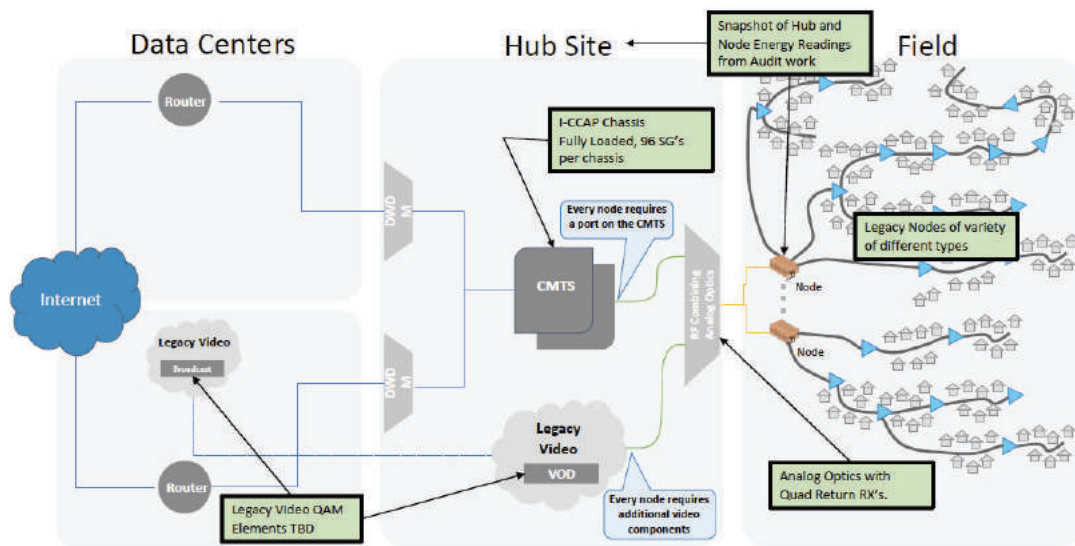


- Collect and analyze any relevant DAA architecture, equipment, and deployment data to model energy usage over time across the Shaw footprint.
- Use the model to prove and/or adjust the hypothesis and answer critical energy usage related questions.
- Utilizing the analysis, work with Shaw to draw conclusions with respect to the energy usage profile related to the Shaw DAA deployment plan.
- Analyze OSP powering capacity and architecture, specifically related to the ability of the power architecture to accommodate DAA evolution, as well as additional IoT, small-cell and other non-HFC connectivity alternatives.

To evaluate DAA impact on the network, the team first developed a model detailing energy usage for the business as usual (BAU) network architecture. It should be noted that the term BAU is used in this paper to refer specifically to Shaw's traditional plant build as currently defined. The team then developed a DAA model, based on Shaw's specific DAA implementation process, to compare against BAU. Conclusions were then drawn from the modeling work.

## 2. Business as Usual Network

The business as usual network has been built and developed over time in a manner consistent with MSO practice. The diagram below provides a view of the BAU network



**Figure - 1 Business as Usual Network Deployment**

Shaw implements a tiered facility structure, with the top tier data center/core facilities located in key regions served and connected by a national mesh ring network. From the regional core sites, rings are used to connect to hub edge facilities in the region. Hub/edge facilities support a load of approximately 50,000 homes passed (HP) with capacity to go to 80,000 HP. The HFC distribution network stars out from the hub/edge facility. In the facility, the company connects nodes to an

integrated CCAP (I-CCAP) device, with standard RF optical devices and connectivity infrastructure in place to support the forward and return bandwidth to/from the nodes.

With respect to Shaw's HFC infrastructure, it is  $N + X$ , with current attention to reduce or maintain node size and cascade lengths. Congestion relief in the downstream is achieved by node-splitting, where the node is physically divided to create two nodes, adding an additional optical path to the head-end, effectively doubling the dedicated downstream capacity to the same area. Upstream congestion relief is via a combination of node splits and conversion to a higher return path bandwidth of 85 MHz, techniques which increase dedicated upstream capacity for the area being served.

The company is committed to evolving fiber deep over time, with the ultimate goal of moving towards DOCSIS 4.0 in preparation for a 10G future. Evolution to  $N + 2$  architecture is the first step to that future [1]. In support of this, greenfield plant and any brownfield upgrades are built to  $N + 2$ .

### **3. Approach to DAA**

As noted in the introduction, an important part of Shaw's evolution to 10G is implementation of DAA in its network. DAA itself is a cable industry initiative, spearheaded by CableLabs, and forms a key part of the cable industry's drive to 10G in the future. Network benefits related to DAA implementation include:

#### **Network Efficiency**

- Increased network capacity
- Better end-of-line signal quality, higher modulation rates, higher bitrates
- Better spectral efficiency, more wavelengths per fiber

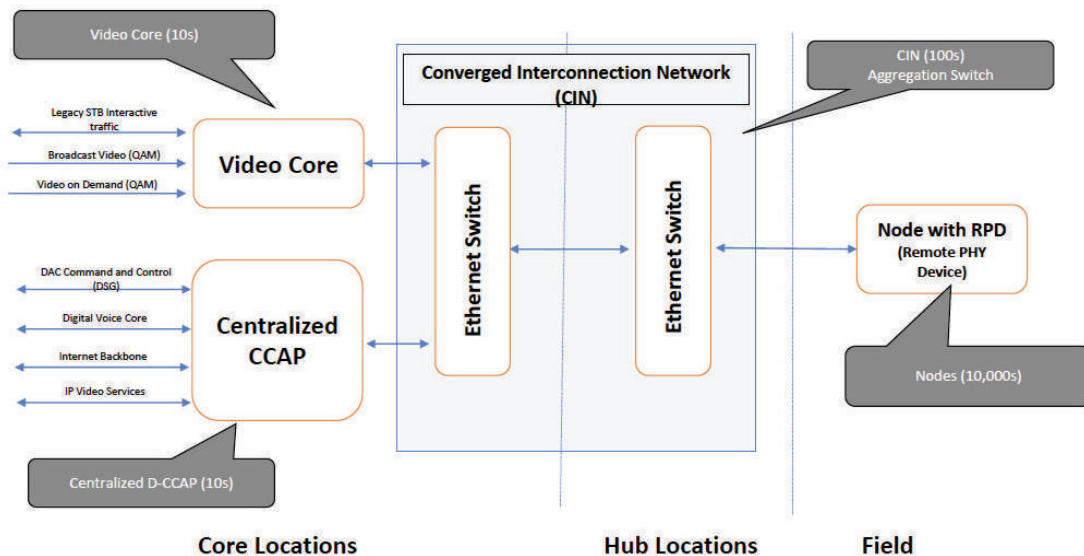
#### **Operational Efficiency**

- Reduced headend power, space and cooling requirements
- Hub consolidation
- Simplified Hub Architecture by leveraging digital optics (eliminates RF combining)

#### **Enabling Convergence**

- Extends IP network into the field
- Ability to leverage common IP Network for multi-tenant applications
- Alignment with FTTx build-out
- Required steppingstone to Virtual CCAP (vCCAP), 10G and the Edge Cloud

A high-level view of the Shaw approach to DAA implementation is shown in Figure 2 below.



**Figure 2 Shaw DAA Implementation**

Key elements of the DAA infrastructure as defined by CableLabs in [4] and to be deployed by Shaw include:

- **Remote Phy Node (RPD):** The RPD is a device located in the node in the network which implements the Remote-PHY specification to provide conversion from digital Ethernet transport to analog RF transport.
- **Converged Interconnection Network (CIN):** The network (generally gigabit Ethernet) that connects a CCAP Core to an RPD.
- **CCAP Core:** A CCAP device that uses MHA v2 protocols to interconnect to an RPD. This device could be a DOCSIS Core, Video Core, OOB Core, RPD Controller (or a combination of these roles). Distributed CCAP (D-CCAP) is the implementation of CCAP Core the Shaw network uses and is used to represent the CCAP core function in this paper.

## 4. DAA Transition

Although Shaw's overall approach to DAA is consistent with that of CableLabs and the industry, a full DAA implementation will be controlled and measured. DAA will be implemented to meet node growth related to greenfield build and node-splits, as well as very selectively in existing brownfield upgrades. To support this evolutionary approach, initially a D-CCAP core and its associated CIN devices are to be placed in Shaw regional core sites. The infrastructure will be used to support the initial DAA node requirement for any hubs served from that core. D-CCAP and CIN network will grow with DAA node need, with D-CCAP chassis' and CIN leaf switches ultimately evolving outward to the hub sites as and when demand in DAA nodes dictates.

With respect to the nodes, Shaw's transition strategy will not entail pro-actively changing existing brownfield nodes to DAA with limited exceptions. Instead, DAA transition will cap legacy node

placement as much as practicable, and use DAA nodes for new greenfield builds, as well as node growth related to node splits. Key attributes of this approach include:

- Focuses on building greenfield plant primarily with DAA nodes.
- In brownfield, implements node splits to relieve congestion primarily in the following manner:
  - Existing node split and new node created with node split will be DAA nodes in the field.
  - Both nodes from the node split will be placed in D-CCAP chassis on 10 Gbps connection via CIN.
  - Existing coax network beyond the nodes stays consistent with traditional HFC node split practices. This is not a change to existing coaxial distribution plant other than to use DAA nodes vs. traditional nodes in the field.

Once a node split has been completed and the new nodes connected to the DAA infrastructure, transition reaches completion with removal of any elements of the BAU I-CCAP infrastructure that has been taken out of service. But as this is an evolutionary process, this happens only after all in-service elements on the legacy devices are completely transitioned to D-CCAP and CIN network:

- I-CCAP chassis removal does not occur until all nodes are removed from the chassis. As the geography of where node splits occur is relatively random and based on many factors, this may be a few years away, given the spread of I-CCAP devices in the network and average number of ports in use on those devices.
- RF transmitters and receivers will be removed and powered down as and when all nodes connected to the chassis have been moved to the CIN. This may happen more quickly, potentially in the first few years after node split, as the number of nodes per chassis is smaller.

At this point, it is not Shaw's intent to actively work to re-arrange BAU node connections to vacate I-CCAP chassis' and/or RF optic modules as utilization on them lessens, but to let natural attrition determine when devices are vacated and can be removed.

Figure -3 below shows detail of Shaw DAA strategy.

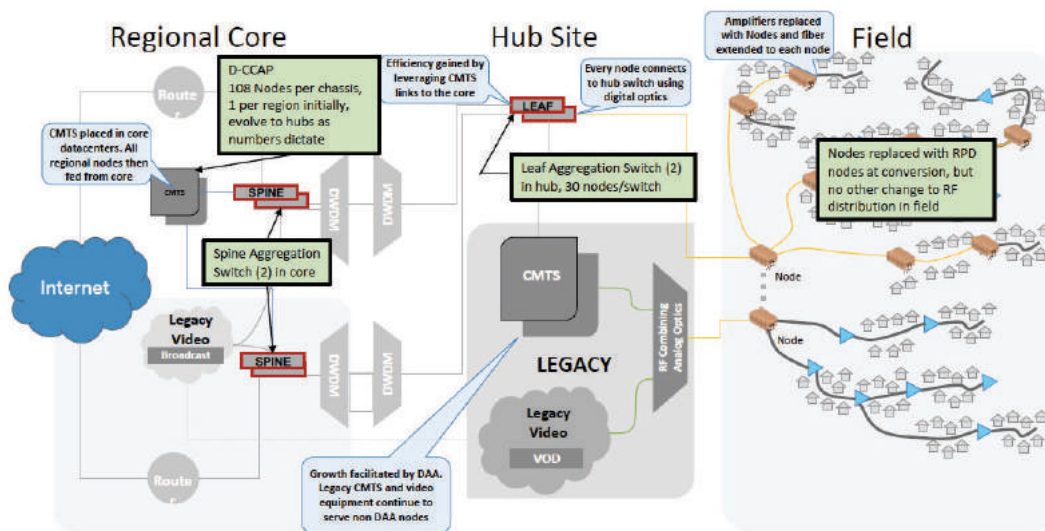


Figure -3 Shaw DAA Implementation

To summarize, with this approach, legacy I-CCAP infrastructure is largely capped in the footprint. The DAA infrastructure operates in parallel to the legacy I-CCAP and RF infrastructure, taking on all growth in nodes due to either greenfield build and/or node splitting. Legacy BAU I-CCAP equipment and supporting RF infrastructure is vacated over time by attrition, with equipment removed as and when service is no longer provided on it.

## 5. Modeling Energy Usage

As with other operators, Shaw's BAU networks grows energy usage year-over-year, due partly to continued growth in greenfield network homes passed, as well as increases in bits transported due to subscriber consumption related to internet use, streaming, etc. Greenfield growth drives additional nodes and actives in the field, along with the associated CMTS and RF optical components to support them. Increasing usage by subscribers creates node splits, along with the CMTS, RF optics, and new nodes required to support them. As subscribers' appetite for data consumption does not appear to be slowing, space as well as energy constraints in existing edge facilities are moving operators towards solutions which solve and/or mitigate these challenges. Evolution to DAA is driven by the needs of the cable operators like Shaw to solve the future space and energy usage problems growth of the BAU network creates if it continues at current pace.


Although in the past, work had been done on energy usage comparing BAU vs. the full implementation of DAA in a facility footprint [5], energy usage comparing BAU to a more evolutionary approach to DAA such as the one used by Shaw has not been modelled and/or analyzed. The following two sections of this paper detail this modeling, as well as results from the modeling work.

### 5.1. Initial Energy Usage Hypothesis

After initial review of Shaw’s BAU scenario and DAA implementation plan detailed above, it was hypothesized that during the transition period from BAU to DAA, initially total energy usage would increase in comparison to BAU, before ultimately decreasing over time. This was due to three key reasons:

- Initially DAA D-CCAP, switch, and router equipment is added in core and hub sites to facilitate DAA, duplicating BAU network resources servicing nodes today.
- Removal of I-CCAP equipment will not occur until an individual chassis is completely vacated. This could take years.
- The DAA nodes in OSP are higher power devices than the BAU nodes they replace.

Figure -4 shows initial high-level modeling of percentage change to BAU associated with DAA implementation and transition done to form the initial hypothesis.



**Figure -4 DAA vs. BAU Percentage Energy Change Initial Hypothesis Modeling**

The initial high-level modeling represented in Figure 5-1 supported this hypothesis, showing an increase in energy usage in the early years related to deploying DAA in this manner. Ultimately, over time as DAA predominates and more BAU equipment is removed, DAA power does reduce, and go lower than BAU. That high-level modeling indicated timing of the cross-over point for DAA producing lower overall energy usage than BAU was a function of how fast DAA was implemented in the network, as well as how quickly I-CCAP infrastructure was removed from the network.

## 5.2. Energy Model Assumptions and Structure

To validate the initial hypothesis, a more detailed model was created. The detailed model starts with a base model for energy use in the BAU network. Layered on top of the BAU base is a model detailing changes related to implementing DAA in the plant, incorporating the Shaw transition assumptions. This allows comparison of BAU vs. DAA specific to the company approach. The base BAU model as

well as the DAA overlay use company supplied and industry data, as well as current company practices for developing and growing the network.

Model Year 0 was specifically designed to align with Shaw current plant and energy information garnered from energy audit work performed in 2018. Key elements included:

- **Network footprint volume:** Used Shaw overall plant data in 2018 for HP, nodes, etc. as a starting point for the base BAU model.
- **OSP power per node:** Calculated OSP power per node using OSP PS data from audit performed in 2018 on calendar year 2017 data, combined with total node numbers from the plant data.
- **Hub/Facility power per node:** Calculated hub facility power per node using data from a sample of stand-alone hubs with known number of nodes connected.

From this data, total OSP and facility power across the Shaw footprint was calculated for Year 0, specifically matching audit data.

Moving in time from Year 0, the model changes homes passed and number of nodes year-on-year into the future using key drivers related to greenfield expansion and node splits. These include:

- **Greenfield expansion:** 1% increase in HP per year. This assumption is based on typical industry data for greenfield in developed markets. This was used to drive node growth due to greenfield build in the model.
- **Percentage node splits in year due to TB growth:** 5% of nodes from previous year. This was based originally on industry data, although a sample of Shaw data supported this assumption as reasonable. This was used to drive node growth in the model due to node splits.

The node growth resulting from the above drivers changes year-by-year energy usage in OSP and hub power based on the following key assumptions:

- **Facility equipment power:** RF optics power, CMTS chassis power, etc. as provided by Shaw for specific devices used in the company network. These numbers were used to calculate energy use changes in critical facilities for node growth related to plant extension and node splits in the year.
- **CMTS chassis utilization:** Model assumes one node per port per Shaw current practice, and 66% port utilization based on experience/industry norms. CMTS chassis' in the model grew as a function of node growth in combination with this utilization assumption.
- **Facility PUE:** Facility PUE was assumed to be 2.0, based generally on conservative industry norm. This is needed in the model as it drives changes for non-equipment energy (i.e. HVAC, lights, etc.) in facilities.
- **Brownfield node power:** Model uses per node power calculated from 2017 PS audit throughout the model.



- **BAU greenfield node power:** As greenfield nodes target  $N + 2$ , in theory they will be smaller than the brownfield assumption around node size and power. As such, greenfield node power is assumed to have W/HP consistent with a brownfield node.
- **BAU node split power:** After a node split, total power for the sum of the two split nodes is incremented by the power difference between the new node and the amp it is assumed to be swapped for.

For the DAA model, many of the base assumptions were similar to BAU. As with BAU, DAA model uses Shaw overall plant HP, nodes, etc. to model a Year 0 starting point similar to BAU today. Key growth drivers around greenfield growth percentage, node split percentage, PUE, and chassis utilization year-on-year are the same, as well.

The DAA model differs largely due to implementation and transition assumptions. Implementation drives changes to per node energy usage in the following ways:

- **OSP power per node:** DAA power per node is adjusted by the increment in power associated with replacement of the standard node with a DAA node.
- **Hub/Facility power per node:** Facility power is adjusted using DAA equipment elements instead of legacy equipment for the added nodes.

The DAA model itself is driven by the Shaw transition strategy. Key elements include:

- **DAA nodes only deployed in greenfield plant:** All plant extension greenfield nodes will be DAA nodes and therefore connected to D-CCAP chassis.
- **Node Splits:**
  - Existing node and the new node created from a split will be DAA nodes in the field.
  - Both nodes from the node split will be placed in D-CCAP chassis on 10 Gbps connection via CIN.
  - Existing node being split will be removed from I-CCAP chassis.
  - RF Optics connected to existing node will be disconnected.
- **Legacy I-CCAP:** I-CCAP infrastructure is capped – all node “adds” due to greenfield build and/or node splits become DAA and are placed on D-CCAP chassis.
- **BAU I-CCAP and RF Optics removal:** These items are removed but only after all BAU services are completely transitioned. Following assumptions are made in the model made for removal timing:
  - RF Optics removed in the same year I-CCAP node is moved to D-CCAP as part of node split.
  - Full I-CCAP chassis’ are removed three years after initial node moved to D-CCAP (i.e. starting year 4 in the model).




- **Evolutionary Strategy:** Strategy does not include targeting hubs and/or CMTS chassis/RF optics for complete conversion, and/or potential transfer of nodes to BAU ports to clear equipment for removal.

## 6. Findings from Modeling

Modelling work focused on a ten-year view of energy usage for each of the two options. Ten years was chosen because it is far enough in the future to see trends, but not beyond the point where the assumptions made could still seem applicable. Output was generated with respect to OSP plant energy usage, hub/facility energy usage, and combined usage, across the whole Shaw footprint. Data and analysis from that work is contained in this section.

### 6.1. BAU

BAU modeling shows future growth in energy usage in the footprint as assumed new greenfield build and node splits drive additional nodes. Figure -5 below shows the changes year-on-year.



**Figure -5 BAU Energy Growth**

BAU OSP growth in the overall footprint is driven by a combination of greenfield expansion, as well as growth due to node split. Key impacting items associated with this were:

- Additional 1% of greenfield plant built in the footprint. This was the primary driver of BAU OSP energy use growth.
- The slightly smaller per node power specifically on the new greenfield nodes as at N + 2 they are smaller and require fewer active elements to be powered.
- The slightly higher power in split nodes due to swap of an amp with a node.


With respect to the individual components, OSP energy grows at a compound annual growth rate (CAGR) of 1.2%, making it 13% higher in Year 10 in comparison to Year 0.

Facility/Hub power grows at a much higher CAGR of 3.0% over 10 years. This is driven almost entirely by the addition of I-CCAP chassis', RF optics and infrastructure needed to support added greenfield nodes, as well as the nodes added via node splits. In comparison to Year 0, Year 10 usage is 34% higher in BAU.

Combining the effects of OSP and facility/hub growth, overall energy growth for BAU implementation had an energy usage CAGR of 1.9%. In comparison to Year 0, Year 10 usage is 20% higher. This combined view forms the basis of the BAU situation for comparison to DAA.

## 6.2. DAA

DAA modeling shows future growth in energy usage as well. As with BAU, new greenfield build and node splits drive this, tempered by the DAA transition drivers. Figure -6 below shows the changes year-on-year.



**Figure -6 DAA Energy Growth**

In relation to the individual components, DAA OSP energy is impacted largely by the same key drivers as BAU. Energy usage CAGR is 1.3% over the ten-year period, making energy usage 14% higher after 10 years. Both figures are slightly higher than BAU due to the swap out of legacy nodes with higher powered DAA nodes.

Facility/hub power growth undergoes much more meaningful changes from DAA implementation. Usage per year is higher in early years (2.1% - 2.2% per year), due to the addition of D-CCAP infrastructure coupled with the delay in removing legacy equipment related to DAA transition strategy, particularly the I-CCAP chassis' equipment. Starting in Year 4, the model begins the process of eliminating legacy I-CCAP chassis', and DAA growth approximately halves. The combination of removing I-CCAP chassis' from the network, and taking up port growth from Year 1 with lower power D-CCAP chassis', works to minimize yearly growth from that time to the 1.0% per year range.

As such, the facility energy use CAGR is 1.4%, with overall growth over the ten-year period slightly above 14%, both much lower than BAU.

The combined effect of the DAA transition strategy yields an energy use CAGR over 10 years of 1.3%, driving growth for the combination of OSP and facility to be 14% higher in Year 10, both lower than BAU. Even with the migratory approach, ultimately the swapping out of I-CCAP for D-CCAP, and swapping RF optics for CIN components, overcomes the added energy for DAA nodes in the OSP, leading to lower energy usage in comparison to BAU. Using the Shaw DAA transition strategy, energy usage will increase in future due to growth in plant and continued node splits – but DAA works to lessen that growth in future in comparison to BAU.

After modelling the Shaw DAA transition strategy, questions arose around how this might compare to energy usage for a complete conversion of a facility/hub to DAA in Year 1, as the company was contemplating this for a hub due to reasons related to an existing facility,. Although previously published work [5] on this topic was written from a more generic perspective, modelling of the Shaw specific use case for implementing DAA immediately in a hub provided some guidance.

Using energy audit data, along with homes passed, node and I-CCAP chassis data for facilities in the Calgary area, year one full DAA was modelled. The model used the following assumptions, yielding the following results


- **OSP DAA immediate implementation:** To calculate this, it was assumed all nodes within each of the hubs were swapped with DAA nodes. This drove a 4% increase in OSP power, solely related to the immediate swap out of lower powered existing nodes with higher powered DAA nodes.
- **Hub/facility DAA immediate implementation:** To calculate this, all I-CCAP chassis' in the hubs were swapped for D-CCAP. Additionally, all optical transmitters and receivers were eliminated, replaced with a CIN network to which the DAA nodes were connected. This resulted in a 42% decrease in hub/facility power, primarily due to the I-CCAP for D-CCAP platform swap.
- **Overall DAA immediate implementation:** The combination of these two effects drives a 5% decrease in overall energy usage. This relatively modest overall percentage reduction is a function of the heavy weighting towards OSP energy in the sample areas, which at ~80% is the majority of the total energy usage. Even though energy savings in the facilities for this approach would be meaningful, because of the large OSP component, they have a more limited impact on overall energy usage.

It should be noted that although not specifically a part of the work, high level estimation was that in addition to the significant savings in facility energy, fully implemented DAA would reduce space to somewhere between a third to a half of what it was pre-DAA.

Although meaningful energy and space savings would be achieved with immediate conversion to DAA in a facility, it is unlikely the savings associated with that change would offset the capital required to accelerate DAA implementation on its own. In special cases, lease challenges and/or space/energy limitations might necessitate such a change. The ability to use an all DAA implementation to meet a particular building and/or facility challenge is an important tool an operator can use. Although there are real benefits in energy and space savings in comparison to BAU to a full DAA implementation on Day 1, other rationale beyond those savings would be required to justify such a conversion.

### 6.3. DAA vs. BAU

Figure -7 compares year-on-year percentage difference in energy use change between BAU and DAA.



**Figure -7 Percentage Difference DAA vs. BAU**

This comparison of DAA energy growth to BAU energy growth shows the Shaw DAA transition strategy, even with its evolutionary approach, delivers energy usage reduction against BAU over time. Overall energy usage stays roughly the same in the early years, with the small decrease in facility energy offset by a slight increase in the much larger OSP component. This changes in the later years, as ultimately removal of I-CCAP components over time accelerate the energy usage reductions of DAA versus BAU.

Learnings from the modeling and subsequent analysis associated with it include:

**Early years impact of dual DAA and BAU infrastructure:** Section 5.1 hypothesized DAA energy use would increase in the early years with placement of dual DAA and legacy infrastructure for a period of time, reducing in comparison to BAU after I-CCAP removal started. In fact, the model showed the impact of this was negligible. Analysis indicated capping the placement of much higher power I-CCAP devices, and adding new capacity needs on newer, lower power D-CCAP devices in the DAA model, mitigated the added power of the dual infrastructure. A learning from this work was that the key to avoiding an early-year increase in energy usage related to this strategy is to cap I-CCAP and put all node growth in facilities on D-CCAP infrastructure.

**Timing and size of DAA reduction – I-CCAP removal assumption:** As expected, the size and timing of the DAA reduction varies with the assumption as to when I-CCAP chassis' would start to be

removed. From a modeling perspective, earlier removal improves DAA in comparison to BAU, later removal lessens the energy improvement of DAA. The assumption of beginning the removal of chassis' in Year 4 equal to the number of ports vacated in Year 1 is an estimate trying to approximate when the nature of node split locations would fully vacate I-CCAP resource over time. If a more exact estimate of this is required, future work could look at adding depth to this part of the modelling based on real world experience, as well as producing quantitative analysis as to impact of varying I-CCAP removal starting point assumption in the model.

**Timing and size of DAA reduction – improvement via consolidation:** As noted in section 5.2, the DAA transition strategy does not include targeting lowly utilized equipment chassis' for complete conversion, to clear for removal. In fact, the timing and size of the DAA savings would be improved with a strategy for setting a minimum utilization point after which legacy architecture would be consolidated. Because a nominal spend would be required to accomplish consolidation, scheduling should be done strategically to minimize costs. One would not want to see this DAA implementation strategy lead to the continued use of a number of rarely-utilized I-CCAP and RF optic chassis' for a long period of time.

**Facility/hub space savings:** Although this work focused on energy usage, it should be noted DAA would lessen space requirements in the facility/hub locations. The DAA transition strategy slow-rolls the I-CCAP to D-CCAP evolution, so space savings will be evolutionary in much the same way energy usage evolves. Although conversion from I-CCAP to D-CCAP platforms yields minimal space savings for the chassis' themselves, elimination of RF optics, as well as the associated RF combining equipment over time, would free up significant space in the facility. The high-level estimate was that, with fully implemented DAA, facility space would be in the order of half to a third of the space used today. Future work specifically related to space would be recommended to better understand and quantify this.

**Future of vCMTS:** DAA and D-CCAP are both part of Shaw's network evolution strategy. Additionally, Shaw is evaluating the potential to evolve the D-CCAP architecture to vCMTS as and when that evolution makes sense. vCMTS holds the promise of housing the CMTS functionality in more space and power efficient white-box storage and server devices, potentially further improving energy and space efficiency in facility/hub devices. Future work to assess energy impact of vCMTS would be recommended as and when evolution to the platform is being planned.

## 7. Outside Plant Power Network Capacity Analysis

As noted in Section 2, Shaw BAU OSP network is standard HFC with ~N+X architecture with plans to maintain or reduce node sizes and cascades lengths. Consistent with HFC deployment norms, the company has implemented a distributed power infrastructure in the OSP to power nodes and active devices in the HFC network. In addition to the nodes and actives, the OSP power infrastructure also supports additional non-HFC devices e.g. small cell, Wi-Fi AP's, etc.

As a part of network evolution of taking fiber deeper, the company intends to use existing power supplies and power supply locations as much as possible. In addition to accommodating network evolution, the company expects the current OSP power infrastructure to be capable of incorporating more of the additional loads related to non-HFC equipment in the future. To date, Shaw has connected in the range of 10,000 Wi-Fi AP's to the coax network, and as an MNO, have connected 4G LTE small cells and are anticipating connecting 5G, both as noted in [3]. The table below outlines the approximate power draw for different types of devices which may be connected to the power network.

**Table -1 Non-HFC Device Typical Power**

| Device            | Power Consumption (per device) | Quantity per Site | Total Power/Site                    | Range                                               |
|-------------------|--------------------------------|-------------------|-------------------------------------|-----------------------------------------------------|
| <b>Small Cell</b> | 60-180W per device,            | 1-3               | 60 - 500W<br>300W-320W/site typical | ~250M (3.5GHz CBRS)                                 |
| <b>Wi-Fi AP</b>   | 15-35W                         | 1                 | 25W                                 | ~200M (2.4GHz)                                      |
| <b>IoT</b>        | 60W                            | 1                 | 60W                                 | 2-4Km                                               |
| <b>Gateway</b>    | Modem (15W-20W) + Load         | 1                 | 80W – 140W                          | For Cameras, etc.<br>- depends on device under load |

Source: Industry Data

To better assess the ability of the OSP power infrastructure to support these goals, as well as understand spare capacity in the power network, the energy audit work performed on the OSP power supplies was used to analyze utilization of the power infrastructure.


The audit found that Shaw's OSP PS population totaled a little over 20,000, of which 95%+ were analyzed as a part of the audit work. With respect to voltage, data indicated ~75% of the PS's were 90V with the vast majority of remaining PS's at 60V, yielding an average plant voltage of 83V. Status monitoring data yielded an average current of 5.8A. Both the 90V/60V split, as well as the ~6A average current, are typical of MSO peer operators [2].

Power supply output voltage and current as acquired via the status monitoring system was used to analyze the PS network. Load investigations were based on Volt-Amperes (VA) calculations from this data, since the output power factor is not known. Key findings from the data included:

- Average output VA across the footprint is 482 VA. This is the average usage for each of the PS's in the network.
- Average capacity for the mix of PS's in the power supply universe is 1286 VA. This implies primarily 15A PS's in the network, with some mix of smaller and larger variants.


This leaves on average ~800VA available across the footprint. As noted in [2], given the majority of the PS's in the footprint are 15A capacity, good practice dictates power supplies target to run close to 80% of load, or with 3A available. This would equate to a minimum spare capacity requirement for standard network of ~250 VA. At an average of 800 VA available capacity, the power supply network appears to have more than enough headroom for future growth.

While averages provide important data, distribution of the count of power supplies around the average can provide operators a more informed view. Distribution of the power supply count by available VA is as shown in Figure -8 below.



**Figure -8 PS Distribution of Remaining Output VA**

As demonstrated above, the vast majority of power supplies have in excess of 200 – 300 VA of capability minimum as headroom recommended. Figure -9 re-structures the data above with respect to the percentage of power supplies in each category.



**Figure -9 Distribution PS Remaining Output VA %'s**

Key takeaways from this analysis include:

- 95%+ Power Supplies have minimum recommended headroom of 250 VA.
- 95% of Power Supplies have 300 – 400 VA or more, providing minimum headroom plus some added capacity for almost all PS's.



- 67% of Power Supplies have 600 – 700 VA or more, leaving minimum 350 VA over recommended minimum for additional loads.
- Average 800 VA of available capacity implies on average 550 VA over recommended minimum for additional loads.

Although the addition of DAA and continued node splits as Shaw evolves to N + 2 architecture will have an effect on reducing headroom, as noted above, it is the company's intent to manage any additional power needs related to network evolution and the addition of non-HFC elements from the existing power supply capacity. This analysis concludes that Shaw's OSP power architecture generally has the capacity to continue supporting the evolution of the HFC network, as well as additional small cell, IoT and gateway devices into the future.

## 8. Conclusion

The journey to 10G will push Shaw's network to evolve. Implementing DAA will be a key element of the network of the future. To properly balance timing and cost with needs and benefits, the company has chosen an evolutionary strategy for DAA implementation. This strategy will embed DAA into the network in a controlled manner, focusing implementation in greenfield nodes, as well as using DAA nodes in node splits at first, and proactively implementing DAA in existing brownfield only in very limited cases.

Even though it is an evolutionary approach, the Shaw DAA migration still provides energy savings over time in comparison to BAU implementation. Modeling work performed on the company's DAA implementation vs. BAU showed the following:

- DAA OSP energy slightly higher than BAU at 14% vs. 13% cumulative ten-year growth, due primarily to higher powered DAA nodes replacing existing nodes.
- DAA facility energy decreasing to 14% cumulative ten-year growth vs. 34% for BAU, due primarily to swap of I-CCAP with D-CCAP, as well as swapping RF optical devices with CIN elements. This is where the savings of the DAA transition are.
- DAA overall increases energy 14% vs. 20% cumulative ten-year growth for BAU. The Shaw evolutionary approach to DAA implementation slows but does not eliminate BAU infrastructure energy growth.

Capping I-CCAP growth and moving node growth to the lower energy DAA platform mitigates the initial power increase of dual D-CCAP and I-CCAP infrastructures. Real energy reduction in comparison to BAU begins when I-CCAP infrastructure is fully vacated and removed from facilities. Timing is impacted by speed of DAA evolution – if the process is accelerated, energy usage benefits will accelerate, and if it is slowed, energy usage savings will similarly slow. Energy usage reduction speed is also impacted by any actions that drive a less random, more geographically targeted DAA implementation. The faster I-CCAP and RF optical equipment can be fully vacated and eliminated, as opposed to lingering in service at low utilization, the faster energy usage benefits will accrue.

The benefits of immediate conversion to DAA were examined. Modeling from Calgary data indicated immediately converting a hub geography to DAA reduces facility power significantly (~42%), but overall impact is more modest (-5%). This is largely due to a large proportion of total energy (~80%) was OSP energy in the sample areas used – other areas where OSP energy was a more typical 70/30



split would see greater benefit. The small OSP increase also played a minor role. Although facility improvement in both energy usage and space reduction would be meaningful in an immediate transition of a hub to DAA, on their own it would be assumed they would not justify the cost. Full DAA implementation in a hub geography, however, could be a useful tool for solving specific space and power challenges, should they arise.

Finally, energy audit data was used to examine headroom in the OSP power infrastructure to handle DAA evolution, as well as additional non-HFC devices intended to be connected to the power infrastructure in the future. The data showed:

- Average headroom adequate to support evolution to DAA and fiber deeper, as well as multiple additional devices.
- ~2/3rds of nodes have headroom for even the largest anticipated single load device to be added and still maintain minimum headroom.
- 97%+ nodes have minimum headroom in place (20% capacity).

Although each case needs to be examined on its own merit, in general, the Shaw OSP power infrastructure has ample headroom to accommodate DAA evolution and the addition of non-HFC elements to the power infrastructure.

## Abbreviations

|         |                                             |
|---------|---------------------------------------------|
| 4G / 5G | fourth generation / fifth generation        |
| BAU     | business as usual                           |
| CAGR    | compound annual growth rate                 |
| CI      | critical infrastructure                     |
| CIN     | converged interconnection network           |
| CMTS    | cable modem termination system              |
| DAA     | distributed access architecture             |
| D-CCAP  | distributed converged cable access platform |
| HFC     | hybrid fiber coax                           |
| HP      | homes passed                                |
| I-CCAP  | integrated converged cable access platform  |
| IoT     | internet of things                          |
| LTE     | long term evolution                         |
| metroE  | metro ethernet                              |
| MNO     | mobile network operator                     |
| MAC/PHY | media access control layer / physical layer |
| MSO     | multi-system operator                       |
| OSP     | outside plant                               |
| PS      | power supply                                |
| PUE     | power utilization effectiveness             |
| QAM     | quadrature amplitude modulation             |
| RF      | radio frequency                             |
| RPD     | remote phy device                           |

|          |                                        |
|----------|----------------------------------------|
| TB       | terabyte                               |
| VA       | volt-amperes                           |
| vCMTS    | virtual cable modem termination system |
| Wi-Fi AP | wi-fi access points                    |

## Bibliography and References

- [1] Nadar Foroughi, “HFC Evolution – The Best Path Forward”, *Proc. Of SCTE Fall Technical Forum*, October 2018, Atlanta.
- [2] Todd Loeffelholz, “Opportunities and Challenges of Implementing Wireless – Small Cell/Wi-Fi/IoT”, *Proc. Of SCTE Fall Technical Forum*, October 2018, Atlanta.
- [3] Jennifer Andreoli-Fang, et al, “Blueprint for Mobile XHaul over DOCSIS – How Low Latency Xhaul (LLX) and Other Technologies Make DOCSIS and Ideal Solution for Mobile Xhaul”, *Proc. of SCTE Fall Technical Forum*, October 2019, New Orleans.
- [4] “Data over Cable Service Interface Specifications DCA – MHA v2 Remote PHY Specification”, CM-SP-R-PHY-I14-200323, CableLabs
- [5] John Ulm, Zoran Maricevic, “Giving HFC a Green Thumb: A Case Study on Access Network and Headend Energy & Space Considerations for Today & Future Architectures”, *Proc of SCTE Fall Technical Forum*, September 2016, Philadelphia
- [6] Rajesh Abbi et al, “Powering the future 10G access networks – An End to End Perspective”, *SCTE Journal of Energy Management, Volume 5, Number 1*

# Optimizing Active Components for Extended Spectrum Networks

A Technical Paper prepared for SCTE•ISBE by

**Chris Day**

Design Center Director  
Analog Devices  
3843 Brickway Blvd, Suite 206  
707 890 8926  
chris.day@analog.com

**Joshua Rose**

Product & Applications Engineer  
Analog Devices  
3843 Brickway Blvd, Suite 206  
707 890 8934  
josh.rose@analog.com

# Table of Contents

| Title                                                                | Page Number |
|----------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                 | 4           |
| 2. The Importance of Bias .....                                      | 5           |
| 3. Introduction to Design of an ES Line Extender .....               | 6           |
| 3.1. Legacy 2-stage Line Extenders.....                              | 6           |
| 3.2. Gain Limitations in Single Stage Amplifier Components .....     | 7           |
| 3.3. Two Stage Design for Additional Gain and Negative Feedback..... | 8           |
| 3.4. Common Packaging, Pin Definitions, and Application Artwork..... | 8           |
| 3.5. Options for Higher Gain in Line Extender Design .....           | 9           |
| 3.6. Variable Attenuators.....                                       | 10          |
| 4. Empirical Modeling Work on Early Devices .....                    | 10          |
| 5. Optimizing Gain, Bias, Tilt, and AGC Allocations .....            | 11          |
| 5.1. Cascade Simulator Concept .....                                 | 11          |
| 5.2. Optimizing Variable Attenuator Placement .....                  | 12          |
| 5.3. Optimizing Gain and Tilt Placement.....                         | 13          |
| 5.4. Optimizing Bias Allocation.....                                 | 15          |
| 5.5. Other Effects .....                                             | 16          |
| 5.6. Simulation Summary .....                                        | 17          |
| 5.7. Resulting LE Performance .....                                  | 17          |
| 5.8. Higher Output Levels .....                                      | 18          |
| 6. Noise Considerations.....                                         | 18          |
| 6.1. Downstream Point of Entry .....                                 | 18          |
| 6.2. Upstream.....                                                   | 19          |
| 7. Conclusion.....                                                   | 19          |
| Abbreviations .....                                                  | 20          |
| Bibliography & References.....                                       | 20          |

## List of Figures

| Title                                                                                     | Page Number |
|-------------------------------------------------------------------------------------------|-------------|
| Figure 1 – Simplified Downstream Line Extender.....                                       | 7           |
| Figure 2 – Single and Two-Stage Cascode Amplifier Component Topologies.....               | 7           |
| Figure 3 – Two-Stage Gain over Temperature .....                                          | 8           |
| Figure 4 – Options for Increasing Downstream ES Gain .....                                | 9           |
| Figure 5 – High Gain 3-Stage Design Options.....                                          | 10          |
| Figure 6 – MER Modeling Example .....                                                     | 11          |
| Figure 7 – Effects of Attenuation Distribution.....                                       | 13          |
| Figure 8 – Example Gain/Tilt Topologies .....                                             | 14          |
| Figure 9 – Attenuator Compensating for Gain Shifts due to Thermals .....                  | 15          |
| Figure 10 – SNR Into Doubler as a Function of Stage 1 and Stage 2 Power Dissipation ..... | 16          |
| Figure 11 – Modeled MER Performance of 45 dB Line Extender .....                          | 18          |

## List of Tables

| <b>Title</b>                                    | <b>Page Number</b> |
|-------------------------------------------------|--------------------|
| Table 1 – Examples of Bias Control Schemes..... | 5                  |

# 1. Introduction

Upgrading current hybrid fiber coax (HFC) networks to extended spectrum (ES) introduces many challenges for operators, equipment providers, and component vendors. Historically upgrades have been accommodated by utilizing improved component performance to develop actives that drop into existing network locations. A drop-in strategy means minimal downtime and reduced labor costs.

Given the familiarity of this strategy it's natural to look at upgrading current actives for ES by leaning on newer component technologies to reduce costs while realizing the goals of increased bidirectional bandwidth. The extension to 1794 MHz might be difficult enough to force a departure from the standard architectures by adding additional active elements at key places in the network. However, before we get to that point the question remains: how far can newer component technology take us?

Major goals for an ES upgrade center around reusing as much of the existing active installations as possible. Ideally, system power supply capacity should not be exceeded, and active upgrading should consist of replacing old equipment trays with new ones without changing locations or re-splicing cables. To not disturb operating levels to existing customers, current output levels are maintained. Any additional energy located in upper bands must be facilitated through the use of newer, more efficient amplifier components powered from the same or lower DC power as today.

Historically, the cable industry has struggled to employ a sufficient workforce during major upgrade cycles. While client-side silicon capable of meeting the full DOCSIS 4.0 requirements may be a way off, it would be advantageous to have ES capable actives soon that are both legacy friendly and easily switched to an ES configuration. ES capable actives can then be deployed as soon as they are ready which will help smooth demand on the workforce when the new modems are ready.

The ES upgrade challenge is like increasing the link budget of a legacy communication system, much of which has been based on older component technology. The incremental losses at 1794 MHz, particularly for customers in unfavorable tap locations with lesser-grade coaxial cable served off long drops can be substantial. Rolling back the quadrature amplitude modulation (QAM) rates for these customers works counter to the purpose of the upgrade. The transmit (Tx) performance from the network active cascade and the receive (Rx) performance at the point of entry (PoE) are key parameters available to offset the detrimental effects of increased losses. In most cases today, components operating in the field provide lower levels of Tx and Rx performance compared to what is readily achievable in a new generation of components using technologies that now serve key locations in wireless networks.

Other improvements and optimizations outside the scope of this paper provide additional performance margin. Low density parity check (LDPC) codes provide distinct advantages over existing Reed Solomon error correction. Similarly, there has been much study about how to best allocate an amplifier's total composite power (TCP) capability by offsetting and adjusting the upper spectrum to 1794 MHz to best fit the physical plant realities.

This paper focuses on how Tx and Rx performance at key locations in the network may be improved to smooth the path for the emerging ES upgrade. In particular, the design methodology for a new set of components targeting optimum Tx power under a fixed DC budget for an ES line extender (LE) is presented. We propose a multiple-stage integrated-circuit-based design approach to optimize overall performance. Likewise, achievable Rx performance is shown as an opportunity for point of entry (PoE) devices.

## 2. The Importance of Bias

Class A amplifiers have served the cable industry well. They provide a good combination of bandwidth and linearity necessary to handle legacy analog TV encoding. Fortunately, networks have been upgraded over time to maintain the high level of linearity needed for analog signals, which serendipitously makes them excellent candidates for handling high level QAM signals in the high bandwidth digital era.

Of course, the downside of Class A stages is their poor efficiency. Achieving high linearity takes a lot of DC bias. Amplifier crash point, the RF output power level where modulation error ratio (MER) rapidly falls as input drive is increased, is likewise dominated by biasing considerations. Although methods of linearization, such as digital pre-distortion (DPD) and legacy analog pre-distortion can provide worthwhile benefit, ultimately crash levels and the ability to increase Tx TCP levels are still constrained by bias. Once a transistor runs out of voltage or current as it traverses along its load line it is no longer capable of reliably carrying information.

It follows then that to maximize the Tx TCP performance attention should be focused on how to optimally allocate bias in the cascading component stages. Higher TCP levels are usually within reach, as demonstrated in recent fiber deep output stages, but at cost of additional DC power that probably eclipses the available legacy power of a high percentage of the networks operating today.

There are any number of ways to adjust bias conditions in Class A amplifiers. The most common and obvious involves adjusting the quiescent current of the stages within a design. However, as equipment vendors seek to differentiate their offerings, it's worthwhile to note that an additional level of efficiency is available by incorporating the supply voltage as part of the design optimization. On this point many vendors have been inflexible on supply voltage and consequently left performance on the table to their competitive disadvantage.

Table 1 highlights the costs and benefits of a few examples of bias control schemes. As ES amplifiers emerge, greater thought can be given to the benefits of incorporating a network management system (NMS) as part of an intelligent network. The techniques below are not mutually exclusive and may be combined to innovate equipment designs. Here the term "bias" is inclusive of both voltage and current being applied to an amplifier component.

**Table 1 – Examples of Bias Control Schemes**

| Technique                           | Description                                        | Advantages                                          | Disadvantages    |
|-------------------------------------|----------------------------------------------------|-----------------------------------------------------|------------------|
| Factory / Field Bias Set            | Adjust once and forget                             | Low cost                                            | Inflexible       |
| Active Bias Control                 | Localized control loop                             | Minimizes variations unit-unit and over temperature | Minor cost adder |
| Remote Bias Adjustment              | NMS monitoring and control                         | Configurable to network differences                 | Added cost       |
| Automatic Bias with Static RF Level | Servo based on required or NMS controlled RF level | Ease of deployment and optimized efficiency         | Added cost       |

| Technique            | Description                                               | Advantages                            | Disadvantages                                                                                           |
|----------------------|-----------------------------------------------------------|---------------------------------------|---------------------------------------------------------------------------------------------------------|
| Envelope Tracking    | Dynamically adjust based on derived signal envelope       | Significant improvement in efficiency | Must know envelope condition and have means to suitably adjust bias conditions, cost of added circuitry |
| Active Linearization | Dynamically adjust based on full spectrum of input signal | Significant improvement in efficiency | Requires much higher speed device processes                                                             |

### 3. Introduction to Design of an ES Line Extender

#### 3.1. Legacy 2-stage Line Extenders


To help with the bias optimization process for ES line extenders, we considered the interaction of gain, corresponding RF output level, and bias expenditure of a hypothetical LE design. For a given RF input level the amount of gain in a stage is intimately connected with its RF output level requirement. We considered how we could develop a family of amplifier components that would optimize cascaded LE performance while fitting within the legacy power available in today’s line extenders. By carefully distributing the amount of gain and bias in intermediate stages we sought to maximize bias available for the final Doubler stage, thereby hoping to maximize the Tx TCP and MER budget.

Of course, line extenders have additional loss elements in the RF path, such as diplexers, automatic gain control (AGC) circuits, directional couplers, and tilt equalizers. Since one goal is to maintain legacy levels in an ES deployment the amount of tilt needed considerably rises. Because most of the tilt is accomplished with passive circuits the cascaded gain must be increased. In addition, the location of this tilt loss must be carefully considered since it places added burden on both the cascaded noise figure and distortion performance of the line extender.

As a starting point, consider a simplified downstream LE block diagram in Figure 1. Two stages of amplification are commonly used in most line extenders today. Depending on how RF output level control is implemented, downstream LE gain for a 1002 MHz design ranges from 39 to 33 dB. An input “Push-Pull” and an output “Doubler” combine to provide approximately 48 dB of amplification. Interstage losses account for the difference between total LE gain and gain provided by these amplifier components.

Cascaded tilt in the range of 8 to 18 dB are commonly configured by applying appropriate tilt modules. Locating most of the tilt directly at the input will deteriorate overall MER performance through thermal noise mechanisms. However, distortion contributions to MER performance will be minimized since each stage operates with a high input tilt. Conversely, locating most tilt between gain stages leads to MER degradation from the driver amplifier since it must drive a higher RF level through the tilt loss.



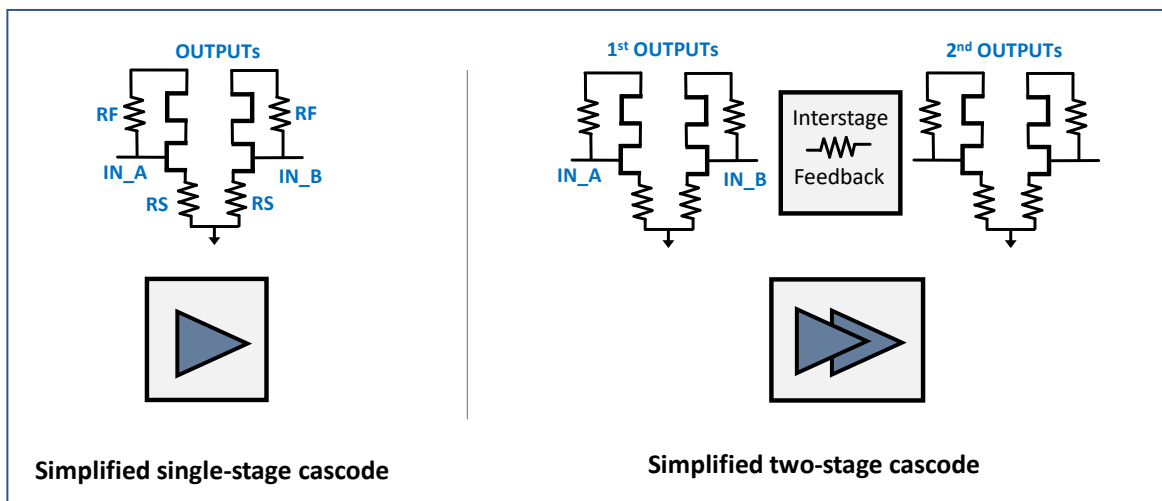


**Figure 1 – Simplified Downstream Line Extender**

To overcome the additional plant loss in an ES buildout, LE gain must be increased by as much as 16 dB compared to legacy 1002 MHz line extenders. Operating tilts will naturally increase to around 22 to 26 dB, meaning that thermal noise contributors to MER degradation will be more difficult to manage. Add to this the desire to carry high levels of QAM in longer RF cascades such as N+4 or N+6, and the problem of where to locate the additional gain and tilt becomes non-trivial.

### 3.2. Gain Limitations in Single Stage Amplifier Components

Most gain stages in cable networks use the familiar cascode topology in a push-pull configuration. Figure 2 shows the familiar single stage cascode topology. We consider a cascode a single stage amplifier since feedback is wrapped around both transistors in the configuration. A combination of series (RS) and shunt feedback (RF) is used to set impedances and manage gain flatness. Due to limitations in intrinsic transistor transconductance and bandwidth, gains of single-stage 1800 MHz amplifiers max out around 23 dB. However, in our testing a slightly lower gain 21.5 dB version has shown superior efficiency performance.



**Figure 2 – Single and Two-Stage Cascode Amplifier Component Topologies**

For intermediate level designs using simple baluns on the input and output, practical gains are closer to 19 dB per stage. These devices commonly take a reduced 5 V to 8 V supply voltage and output a suitable level to power moderate amounts of attenuation and tilt leading to the output Doubler stage. Because these designs have lower power consumption, they can be fabricated on a single die leading to more

consistent performance. If standard baluns are used for unbalanced-to-balanced conversions designers can achieve good impedance consistency using little board space. Standard GaAs processes are commercially available to fabricate these circuits leading to acceptable overall cost.

### 3.3. Two Stage Design for Additional Gain and Negative Feedback

Figure 2 also shows a two-stage cascaded cascode amplifier component concept. The additional stage provides a boost in available gain useful in any number of additional series or shunt on-die feedback arrangements. This provides the integrated circuit designer freedom in setting impedance over a wide range of gain levels. It's possible to design a two-stage integrated circuit where fabricating a range of gains is just a matter of changing on-die resistor values. Practical gains for this two-stage approach range from 22 dB to 30 dB.

An example of a single-die two-stage 25 dB gain design with moderate interstage feedback is shown in Figure 3. The high degree of feedback provides consistent performance over temperature and fabrication process variations. Gain stability over temperature for the 25 dB prototype inclusive of input and output baluns is shown below. Bias current variation over temperature was  $\pm 1.5\%$ . Data was taken from a heat sink temperature of -30 deg C, 25 deg C, and 100 deg C.




Figure 3 – Two-Stage Gain over Temperature

### 3.4. Common Packaging, Pin Definitions, and Application Artwork


In many cases developers of equipment need flexibility to adjust their designs with different gains and power levels to serve the wide range of operator requirements. In the past component vendors relied on the SOT115 package to offer different gain levels to the market. By leveraging this approach, we can

design single stage and two stage integrated circuits with a common application circuit and layout. A bill of materials (BOM) change is all that is needed to adjust gain and DC power consumption levels.

For example, an industry standard 5mm x 5mm QFN package with an exposed backside paddle can be used to package a family of intermediate level amplifier components, with gains ranging from as low as 12 dB through 30 dB. With good thermal precautions, power consumptions of up to 4 W can be managed easily, although as will be seen later our initial integrated-circuit designs consumed between 1.7 W and 3.0 W.

### 3.5. Options for Higher Gain in Line Extender Design

With available gain levels from components in mind, a few options for implementing higher gain can be considered. The question becomes how to optimally achieve the higher gain with additional stages in a modified LE block diagram while holding to the power envelope. A few options are shown in Figure 4.




**Figure 4 – Options for Increasing Downstream ES Gain**

A 3-stage design utilizes push-pull and mid-gain stages operating from the lower supply rail which provide up to 19 dB of gain. Tilt and gain control can be accomplished in any number of ways with variable loss circuits to optimize dynamic range. An output Doubler provides up to 21.5 dB of gain. Although simple enough, the 3-stage design will struggle to provide enough gain to serve the entire range of LE gains operators need to drive a wide range of physical plants currently in service.

A 4-stage extension provides good gain, allowing design margin to locate interstage attenuation and tilt networks for best station performance. However, the added cost of the 4<sup>th</sup> stage is a disadvantage considering the anticipated cost pressures on LE upgrade modules.

An alternative is found in a 4-stage design, implemented in 3 amplifier components, termed a 3a stage design in Figure 5. Here the middle stage is a single die design encompassing 2 cascode stages with feedback wrapped within the 2 cascode stages.



**Figure 5 – High Gain 3-Stage Design Options**

### 3.6. Variable Attenuators


Gain adjustment can be performed using variable attenuator circuits around the various stages. Since gain of the line extender can shift from unit to unit and over temperature, variable attenuation is commonly used to adjust for these changes. Maintaining a flat frequency response and consistent return loss over a moderate range of attenuation is of primary importance to reduce cascaded frequency response ripple. Unfortunately, in attenuator design there is often a tradeoff between impedance consistency over attenuation, attenuation range, and minimum attenuation. Traditional approaches have considerable insertion loss while newer designs have less. Because there are a wide range of implementations used by equipment vendors it behooves manufacturers of amplifier components to provide a wide range of gains in a common package and footprint to provide equipment designers the flexibility necessary to reach their goals.

## 4. Empirical Modeling Work on Early Devices

Given the interplay between gain, bias consumption, and output level capability, we sought to construct a measurement-based representation of early prototype 1800 MHz capable devices for use in LE system simulations. We fabricated standard cascode topology push-pull amplifiers on a single gallium arsenide (GaAs) pseudomorphic high electron mobility transistor (pHEMT) die with series and shunt feedback resistors chosen to achieve close to the high end of realizable gain comfortably beyond 1800 MHz. Single-stage and two-stage designs were fabricated. With a single-stage 19 dB gain intermediate level stage as a starting point, we characterized the relationship between bias current and voltage, MER, output TCP level, and various tilts. Our test source exercised 384 to 1794 MHz with a 6 dB offset above 1026 MHz.

Tilt, voltage, and current were exercised extensively to produce a comprehensive data set for these devices. We then developed an interpolative software model for this device which could be used for

continuous simulation of gain and linearity within the bias and tilt space of the original data set. The data collection was thorough enough that the model very closely matches actual device performance. Figure 6 shows an example of output from the device model over a range of bias and output conditions.



**Figure 6 – MER Modeling Example**

We also fabricated early devices for building output Doubler stages. We intentionally kept the question of supply voltage(s) open, knowing that this choice has major ramifications on overall TCP performance and efficiency outcomes. While historically a 24 V rail has been a strict requirement on Doubler output stage designs, the evolution of device technology has opened possibilities for improved Doubler performance using non-24 V rails. Recent 34 V Doublers serving the higher TCP Fiber Deep architecture are one example of what can be achieved with a wider design window. Furthermore, we assumed that the optimum LE TCP efficiency would come from two distinct supply rails – a higher voltage rail for the output Doubler and a lower rail for all other actives in the LE, including upstream gain.

Using the same approach as the 19 dB gain stage, we characterized the 21.5 dB gain Doubler stage and developed a model for use in system simulations.

## 5. Optimizing Gain, Bias, Tilt, and AGC Allocations

### 5.1. Cascade Simulator Concept

Equipped with models for extended spectrum active devices, we set out to develop a model for a complete LE. The intent was to determine the guiding principles for choosing the optimal topology for an LE, including the distribution of gain, tilt, biasing, and placement of any AGC element(s) or other lossy elements. This would both inform the development of a more comprehensive portfolio of extended spectrum active devices and allow us to provide data-supported recommendations for LE layout and

design. In order to extend our active device models into a full LE simulator, we had to develop several additional elements. First, simple passive element models (diplexers, equalizers, plug-in attenuators, and variable attenuators) were developed. Second, bounding conditions and assumptions on passive elements for LE design had to be determined. Lastly, a tool capable of cascading the system had to be developed to work with these models. In order to operate with the non-standard device models, we had developed for the active devices, we chose to create a custom cascade simulator for this purpose.

Starting from a defined un-tilted input and a perfect SNR, the simulation tool operates on an arbitrary topology of LE passive and active elements, cascading MER, noise, and signal in the downstream path. The cascade technique assumes that the MER can be treated as equivalent to uncorrelated noise power for the purposes of cascading. We found this relationship can be practically demonstrated on a test bench. This model does not consider BER; however, the output levels of the active devices in a typical line extender should keep bit errors from becoming a concern for purposes of these simulations. Further, the use of higher modulation orders in OFDM carriers in DOCSIS 3.1 / 4.0 deployments results in BER which is highly correlated to MER and generally not a function of device crash behavior.


We largely focused on simulations of 3-stage topologies, as it quickly became apparent that by leveraging higher gain intermediate stages as outlined above, four active stages should not be necessary except perhaps in the most extreme high-gain, high-loss cases. Even with higher gain early stages, two-stage designs are impractical for all but the highest gain actives and lowest gain requirements for LE modules.

Except where otherwise noted, most simulations were performed for an overall LE gain window of 40-50 dB with the following constraints and assumptions:

- 70 dBmV total composite power at LE output
- > 45 dB MER performance of full LE
- 26 dB output tilt from 54-1794 MHz (24 dB passive tilt, 2 dB from active components)
- 0 dB input tilt
- 2.5 dB loss at both input and output for diplex filters and other passive elements
- 5 dB (or alternatively 3 dB) minimum insertion loss for a well-matched variable attenuator circuit
- Variable attenuator nominally set 6 dB above minimum to accommodate up to 5 dB gain loss for extreme temperature excursions (+1 dB for safety)
- 1.5 dB minimum insertion loss for equalizer/tilt modules
- ~18.5 W total DC power for downstream amplifiers

## **5.2. Optimizing Variable Attenuator Placement**

The first question that we addressed was how to distribute attenuation – namely the placement of variable attenuator(s) as part of an AGC scheme. These could in theory be placed anywhere in the chain, but the practical locations would be between the first and second stage, between the second and third stage, or both with the attenuation split in some manner. The simulation was conclusive on this point – as long as other components are chosen and located properly, the optimal choice is to place any variable attenuation between stages 1 and 2. This keeps the output of the second stage from being overtaxed, and any noise floor concerns are better alleviated by moving a portion of the tilt after the second stage instead of shifting attenuation. This also lowers the overall gain requirement from the actives by a reasonable amount (compared to a split attenuator design). Consider Figure 7, which shows three optimized topologies for a 45 dB gain line extender, with the attenuation distributed as outlined above. Note also that the “Evenly Split” configuration additionally requires several dB more gain from the actives.



**Figure 7 – Effects of Attenuation Distribution**

This result fits with expectations: regardless of the distribution of attenuation, the output from the first stage remains the same. The only device linearity concern that can be addressed is the second stage's output power requirement. Optimally 100% of the attenuation and loss would be between stage 1 and stage 2 for this reason, but noise-related effects preclude this, so we seek to strike a balance by placing all of the attenuation and the “right” amount of tilt between stage 1 and stage 2.

We found that a useful metric for comparing various configurations was to look at plots of the MER/SNR coming in to the Doubler versus the amount of tilt placed between stage 1 and stage 2. The remainder of tilt is then implied to be between stage 2 and stage 3. Removing the Doubler linearity from the equation (though not its gain, since that informs the gain required from the rest of the chain) allows for the contributions of the various configurations to be more easily distinguished. Plotting MER against the tilt allocation provides a useful way to consider two related dimensions of the problem at once.

### 5.3. Optimizing Gain and Tilt Placement


We were now able to investigate questions about optimal gain and tilt distribution in the LE. Consider Figure 8, which shows a variety of possible gain and tilt distributions, grouped into colors based on which of the first two stages has higher gain. . Two Doubler gain scenarios are considered. One scenario uses a Doubler gain of 21.5 dB and a second uses 28 dB gain. For each Doubler gain scenario we consider the effect of allocating the remaining required gain to the 1<sup>st</sup> and 2<sup>nd</sup> stages and use the tilt allocation between stages as the driving variable in the analysis.

The first and most apparent result is the observation that very high Doubler gain is ideal from a linearity perspective. A Doubler with 28 dB gain would enable the rest of the LE to be built with low gain stages without fear of non-linear contributions from the intermediate stage, even with low bias on that

component. With a 28 dB gain Doubler the remaining gain can be evenly split between 1<sup>st</sup> and 2<sup>nd</sup> stages, resulting in the blue plot in Figure 8.

Unfortunately, when designing within the legacy DC envelope for line extenders, process limitations constrain the gain to the low 20s. With a reasonable increase in bias voltage and current for the Doubler stage it would be possible to produce a 28 dB gain output stage with similar linearity to existing 12-14 W, 21.5 dB gain designs. That is generally outside the scope of this paper, but it bears consideration for situations where such a design change is possible.

Considering the scenario where the Doubler gain is 21.5 dB, the analysis shows best MER performance with the 2<sup>nd</sup> stage gain greater than the 1<sup>st</sup> stage. The various curves account for different levels of gain allocations between the 1<sup>st</sup> and 2<sup>nd</sup> stages. Regardless of the split it's favorable to have more gain in the 2<sup>nd</sup> stage with tilts evenly split between stages, resulting in MER into the Doubler above 53 dB. Placing too much gain in the 1<sup>st</sup> stage leads to distortion contributions from the 1<sup>st</sup> stage.



**Figure 8 – Example Gain/Tilt Topologies**

Starting from the known capabilities of the output Doubler, the ideal tilt split and gain of the first two stages can be solved simultaneously. Not surprisingly, given the number of factors which can contribute to degrading the SNR in the line extender, there is no simple set of linear equations which can be derived to produce the correct answer for all conditions. Further, such an answer would remain only a guideline as there are not enough varied active or passive elements available for actual design.

While a direct invocation of the simulation would be necessary to find the exact optimal topology for a given design target, there are two simplified approaches which both give reasonable approximations of optimum LE performance for LE gains in the range of 40-50 dB, given output TCP/tilt and passive loss criteria roughly similar to those outlined above. All gains are at 1.8 GHz:



- 1) Hold the gain of the second stage (~25 dB) and tilt split (~14/10) constant and allow stage 1 gain to vary
- 2) Hold the gain of the first stage constant (~19 dB) and allow stage 2 gain and tilt split to vary:

In general, the second approach proves to be more practical. It broadens the window for optimally splitting the tilt and for getting the correct gain. What this implies is that when the device gain varies due to temperature and the attenuator needs to be adjusted to compensate, the shift in performance as the arrangement of gain and tilt becomes “non-optimal” will be minimized. To demonstrate this, consider Figure 9 which shows an estimation of amplifier gain (and passive loss) shifting as a function of increasing or decreasing temperature, and the result of adjusting the attenuator to maintain 45 dB overall LE gain. The gains in this example are 19 dB for stage 1, 25 dB for stage 2, and 21.5 dB for stage 3.





Figure 9 – Attenuator Compensating for Gain Shifts due to Thermals

#### 5.4. Optimizing Bias Allocation

All the data so far presume some typical biasing for the components in question. Ideally, the overall DC budget for the downstream section of the LE will provide as much DC power as possible to the output Doubler while keeping the first two stages operating at just above the level where they would contribute meaningfully to the overall output SNR. To look at this, we'll consider the second method outlined previously, where stage 1 gain is held constant and stage 2 gain and tilt split are varied. Under these conditions, can any bias be borrowed for the Doubler?



**Figure 10 – SNR Into Doubler as a Function of Stage 1 and Stage 2 Power Dissipation**

Simulation of optimal gain and tilt distribution for a 45 dB gain LE over stage 1 and stage 2 bias can be seen in Figure 10. With the modeled components, it would be possible to extract ~1 W from the second stage bias and transfer it to the Doubler, presuming that a change of incoming SNR to the Doubler from ~54 dB to ~52 dB would not negatively impact output performance. In our model, increasing the Doubler bias by 1 W was break even with the decreased incoming SNR to that stage, because the Doubler was already running at an optimized bias condition.

Optimizing the bias allocations within a given DC budget can be a sensitive process and would be best performed on prototype boards to directly observe the crossover point where improvement in Doubler linearity no longer overcomes degradation in earlier stage linearity.

It's worth noting that in theory, the bias and gain configurations could be simultaneously optimized to produce a slightly more favorable result. The devices were nominally biased well already, and the model would need to be significantly more precise throughout to benefit from such a small optimization.

## 5.5. Other Effects

There are a few other effects which can alter the model that are worth mentioning, the first of which is noise figure. The influence of noise figure for the input device is understood, and while it's relatively small, it grows larger as the overall LE gain increases or if too much tilt is placed after the first stage amplifier. For example, an optimally constructed 48 dB overall gain LE will see a ~0.5 dB change in MER coming into the Doubler for a 1 dB change in stage 1 noise figure. When considering a 42 dB gain LE, the change is only ~0.2 dB. More surprising is the fact that the second stage noise figure should also not be ignored. In this case the discrepancy between low and high gain LEs is more stark – at 42 dB overall gain, the effect of stage 2 NF is less than 0.05 dB, while at 48 dB overall gain, the effect is roughly 0.5 dB MER per 1 dB NF. This occurs because the optimal distribution of tilt and attenuation

pushes the input signal to the second stage as low as possible to avoid non-linearities behavior in that stage.

One last effect worth mentioning is that of the minimum insertion loss of the variable attenuator. We have made a very conservative assumption that this would be around 5 dB, but it is certainly possible with careful design to craft an attenuator with lower insertion loss (IL). In doing simulations with a minimum IL of 3 dB, we found that while it is overall a more favorable configuration, it does not provide as much benefit as one might expect. Less attenuation required between the first two stages allows for more tilt to be placed before the second stage and enables the use of a lower gain input stage (or a lower gain second stage, although this approach is less beneficial per our analysis). Excepting some fringe cases, we found that the overall MER improvement into the Doubler from a 2 dB change in attenuator minimum loss was on average about 1 dB. A lower IL for the attenuator does enable more design flexibility, provided it remains well-behaved through its full attenuation range and is useful for incrementally improving linearity.

## **5.6. Simulation Summary**

In the end, it's impossible to provide hard and fast rules for overcoming all the challenges that designing a line extender entails. However, some guidelines have been provided for managing the interrelated set of variables that must be considered. We hope that these will prove useful in the design of line extenders for 1.8 GHz going forward. The simulations helped us identify the concept of developing a family of interchangeable downstream amplifier components.

## **5.7. Resulting LE Performance**

Using the measurement-based models we considered overall cascaded performance as a function of the tilt spilt using the previously described parameters. The results are shown in Figure 11. Variable attenuation is placed between stage 1 and stage 2. Gains for stage 1, stage 2, and stage 3 are 19.0 dB, 25.0 dB, and 21.5 dB respectively. Coaxial port output TCP is set to 70.0 dBmV and downstream consumption in amplifier components is limited to 18.5 W. Upstream amplifier component prototypes capable of 40 dB gain and 69 dBmV TCP from 108 MHz through 684 MHz consumed 3.5 W. This brings the combined DC consumption to 22 W. The results show relatively consistent MER performance > 45 dB with the best option being an even splitting of the passive tilt.




Figure 11 – Modeled MER Performance of 45 dB Line Extender

## 5.8. Higher Output Levels

Although the focus has been on staying within the existing LE power budget, it is feasible to increase the Tx TCP from the Doubler by adjusting the design and biasing scheme. Prototype Doublers have shown ability to output > 77 dBmV TCP but require near the maximum SOT115 power dissipation of 18 Watts. Getting higher output levels would best be approached with a new packaging design for better thermal and RF characteristics.

## 6. Noise Considerations

### 6.1. Downstream Point of Entry

Thus far we have focused on optimum design for maximizing LE Tx power and MER performance within legacy boundary conditions. The recent DOCSIS 4.0 specification establishes an input power range of +15 to -30 dBmV per 6 MHz channel. Low input levels are likely for distant reaches in a coaxial plant due to a combination of lower-grade coaxial cable, unfavorable tap location, and long drop length. In these cases, the overall SNR performance will be dominated by the Rx sensitivity and not distribution plant MER contributions.

To address the situation, we designed a low noise amplifier (LNA) stage with an integrated low loss linear bypass switch suitable to handle the wide input dynamic range. The bypass mode can be activated about midway within the range thereby alleviating the need for the LNA to have excellent linearity over the full range of input powers.

We used a linear pHEMT process with good noise characteristics. The process is consistent with a low-cost high-volume environment, not unlike mobile device applications where higher performance GaAs processes provide favorable cost-benefit to the network.

Simulated noise figure for the LNA path was <1.7 dB with a forward gain of 17 dB. Considering a secondary gain stage with noise figure of 5.0 dB, the overall noise figure can be < 2.0 dB. Comparing to today's cascaded noise figure of 4.5 dB, it's possible to improve the SNR performance, and hence the achievable data rate, for some locations at extreme locations in the plant. Just as in the case of maximizing Tx performance for added link budget with careful design, improving receiver sensitivity provides similar opportunity. Although improvement in distribution plant MER provides benefit across the serving area, input noise figure improvements at the point of entry are particularly helpful at extreme network locations. Improved receiver sensitivity is available to help overcome the significant link budget implied with Extended Spectrum deployments over lossy legacy infrastructure.

## **6.2. Upstream**

PHEMT devices feature high intrinsic transconductance and excellent bandwidth. However, compared to other process technologies they unfortunately have poor 1/f noise characteristics. Often, 1/f noise corners for pHEMT devices can be around 30 MHz. The 1/f noise corner depends on semiconductor start material properties which are generally not well controlled by epi vendors and fabricators alike who commonly target much higher volumes not sensitive to 1/f considerations. For these reasons pHEMT devices make inconsistent input stage devices in upstream applications. They may make suitable output sections for upstream applications since noise performance is dominated by input stage contributors. On the other hand, bipolar devices have 1/f noise corners well below 100 KHz and are better choices for input stages in upstream amplifiers.

## **7. Conclusion**

Migrating to an extended spectrum architecture will require careful network planning and design of network elements. Given the magnitude of overcoming the substantial increase in passive losses we sought to design a family of components that could best serve the industry in this task. As with past upgrade cycles, newer component technology is available to help facilitate an upgrade with minimal disruption to existing networks. These technologies are best utilized by opening the design window to optimally allocate gain, bias, attenuation, and tilt in ES equipment design.

# Abbreviations

|        |                                                  |
|--------|--------------------------------------------------|
| AGC    | automatic gain control                           |
| BOM    | bill of materials                                |
| C      | Centigrade                                       |
| dB     | decibel                                          |
| dBmV   | decibel relative to 1 millivolt                  |
| DC     | direct current                                   |
| DOCSIS | Data-Over-Cable Service Interface Specifications |
| DPD    | digital pre-distortion                           |
| ES     | extended spectrum                                |
| GaAs   | gallium arsenide                                 |
| GHz    | gigahertz                                        |
| HFC    | hybrid fiber coax                                |
| IL     | insertion loss                                   |
| kHz    | kilohertz                                        |
| LE     | line extender                                    |
| LNA    | low noise amplifier                              |
| LPDC   | low density parity check                         |
| MER    | modulation error ratio                           |
| MHz    | megahertz                                        |
| NF     | noise figure                                     |
| NMS    | network management system                        |
| OFDM   | orthogonal frequency-division multiplexing       |
| QAM    | quadrature amplitude modulation                  |
| QFN    | quad-flat no leads                               |
| pHEMT  | pseudomorphic high-electron-mobility transistor  |
| PoE    | point of entry                                   |
| RF     | radio frequency                                  |
| Rx     | receive                                          |
| SCTE   | Society of Cable Telecommunications Engineers    |
| SNR    | signal-to-noise ratio                            |
| TCP    | total composite power                            |
| Tx     | transmit                                         |
| V      | volt                                             |
| W      | watt                                             |

## Bibliography & References

CM-SP-PHYv4.0-I01-190815 *Data-Over-Cable Service Interface Specifications, DOCSIS® 4.0, Physical Layer Specification*; Cable Television Laboratories

# **A Taxonomy of Fraud Experienced by Network Service Providers**

## **So Many Ways People Try to Take Others' Money!**

A Technical Paper prepared for SCTE•ISBE by

**Kevin Taylor**

Fellow  
Comcast  
Englewood, CO  
Kevin\_Taylor2@comcast.com

**Steve Goeringer**

Distinguished Technologist  
CableLabs  
Louisville, CO  
s.goeringer@cablelabs.com

**Eric Winter**

Assistant Vice President Investigations and Technical Risk  
Cox Communications  
Atlanta, GA  
Eric.Winter@coxinc.com

**Michael Khalilian**

Senior Director  
Comcast  
Philadelphia, PA  
Michael\_Khalilian@comcast.com

# Table of Contents

| <b>Title</b>                                            | <b>Page Number</b> |
|---------------------------------------------------------|--------------------|
| 1. Introduction.....                                    | 3                  |
| 2. Example Fraud Scenarios.....                         | 3                  |
| 2.1. Fraud Example 1 – Account-Based Fraud.....         | 3                  |
| 2.2. Fraud Example 2 – Credential Based Fraud.....      | 4                  |
| 2.3. Fraud Example 3 – Brand-Enabled Fraud.....         | 4                  |
| 3. The Business Context.....                            | 4                  |
| 4. Fraud vs Cyber Security.....                         | 5                  |
| 5. A Fraud Framework.....                               | 6                  |
| 6. Details Relevant for Explaining Specific Frauds..... | 7                  |
| 7. What Use is a Taxonomy? .....                        | 9                  |
| 8. Conclusion.....                                      | 10                 |
| Abbreviations.....                                      | 10                 |
| Bibliography & References .....                         | 10                 |

## List of Figures

| <b>Title</b>                                                           | <b>Page Number</b> |
|------------------------------------------------------------------------|--------------------|
| Figure 1- Fraud in the Business Context.....                           | 5                  |
| Figure 2 - A Fraud Framework.....                                      | 6                  |
| Figure 3 - A taxonomy example of the Online Contact Channel Fraud..... | 9                  |



# 1. Introduction

All service providers and retail businesses experience crime in the form of fraud. Fraudsters seek to monetize cybercrime and other crime by targeting our companies, our partners, and our customers. Annual fraud losses in the US represent a staggering \$170B across all fraud types. Many of these fraud types apply to the provider's technical and financial operations including card fraud(\$32B), online fraud(\$26B), identity fraud(\$16B), check fraud(\$7B), synthetic fraud(\$6B), account takeover(\$5B), and new account fraud(\$3B)[1]. The methods to execute fraud vary from simple to very technical. In the fraud environment operators face today, fraud incidents are often preceded by cybersecurity events. This requires a greater coordination and collaboration between fraud teams and cybersecurity teams. The fraudsters are creative and coordinated – they attack all aspects of our businesses, including care channels, video, wireless, Internet, mobile and voice services. This paper provides an overview of some examples of fraud experienced by cable operators, and then proposes a framework for fraud and a taxonomy, in order for it to be effectively described and discussed.

What is fraud? Why is it important? Operators are seeing an increase in external fraud resulting in service theft, device theft, and brand damage. In some cases, bad actors use compromised personal and payment information to create fraudulent accounts, or they may use Internet fraud resources to create fully synthetic identities. They may use modified devices to provide video and Internet services to customers they are servicing. Other approaches have been identified as well. In some cases, the subscribers working with the bad actor believe they are working with their actual carrier. Both wireline and wireless operators are being targeted across the globe.

There are many types of fraud being executed against cable operators and their partners. These may target the operators themselves, their customers, or even their business partners. We're focused on fraud that impacts the operator's service and the consumers of those services.

Different industry segments and their products experience fraud differently. Voice, internet, mobile, email, and video fraud all have characteristics that are unique to the operator and the product line. Supply chain and fulfillment may be part of fraud, as well – such as when a fraudster arranges for cell phones to be shipped to a certain location or sets up a new cable account resulting in shipment of cable modems and set top boxes.

## 2. Example Fraud Scenarios

This section examines three scenarios identified by service providers as being fairly common. Fraud follows the business context. It will be aligned and in context of the service or product being offered.

### 2.1. Fraud Example 1 – Account-Based Fraud

New accounts can be fraudulently created. These accounts are usually created to steal service. In these cases, the fraudsters often offer a product and service at a substantially reduced price. The fraudster will attempt to set up many fraudulent accounts to service their customers. These fraudulent accounts are set up with either stolen customer information or synthetic identities created for this purpose and to perform other fraudulent activities. In some cases, synthetic identities are set up and managed for years, for the sole purpose of defrauding companies over the lifetime of the synthetic identity.

## **2.2. Fraud Example 2 – Credential Based Fraud**

A second common fraud leverages existing customer credentials. A common scenario starts with “credential stuffing” [2] and ends up as fraud against the customer or service provider. A customer’s credentials are leaked from one of the “credential spills” that happen every month across the globe. These credentials are picked up by “credential testers,” who will test the credentials against many websites, including those of the operator. If the credential is tested as valid on the operators’ network, the “credential tester” will then post the credential for sale on a darknet marketplace. The credential is then bought by someone intent on committing fraud against the operator. The fraudster will use the credential to access the account and defraud the operator. This may be by ordering additional devices (mobile phones, set-tops, or cable modems). In the whitepaper “The Economy of Credential Stuffing Attacks” by the Insikt Group and available at Recorded Futures [3], the authors note that with an investment of \$550 a “credential tester” can make up to \$19,000 doing just the testing and sales of credentials in a matter of several weeks to months. A very lucrative investment indeed! This example helps illustrate the importance of collaboration between cybersecurity and fraud teams in deterring both the testing of credential as well as the fraud enabled by the validated credentials.

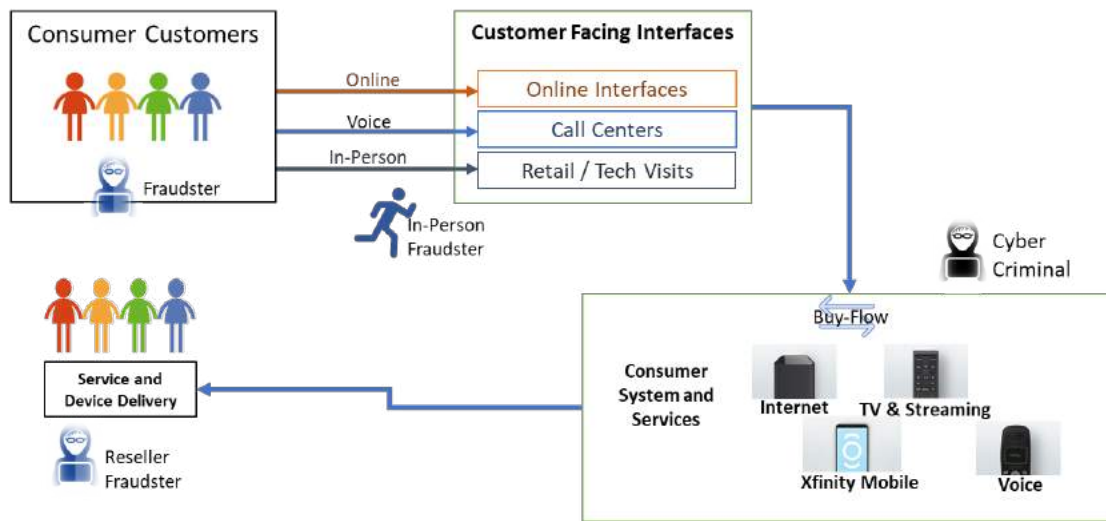
## **2.3. Fraud Example 3 – Brand-Enabled Fraud**

Another example of fraud is one committed using the operator’s brand to defraud the customer. An example of this is a fraudster who calls the operator’s customer, poses as an operator representative and touts a special offer between the operator and a selected brand of gift cards. The fraudster will inform the customer that if they purchase gift cards and call back, they will reduce the customer’s monthly bill substantially. Of course, there is no such offer. If the customer falls for the scam they will purchase the gift cards, call back a number they are provided, provide the information to the fraudster and the fraudster will claim that their bill will be lowered. When the actual bill from the operator arrives, there is no change. The customer calls the operator and then discovers the fraud, which frustrates the customer and damages the operator’s brand.

## **3. The Business Context**

There are many channels through which these frauds are conducted, which are depicted in Figure 1. Fraudsters may access the customer’s operator web portal. They may call in and work with the interactive voice response (IVR) service or a customer care agent and escalate the call to full social engineering through direct, in-person interaction with retail employees, at kiosks or stores. On the other side, fraudsters may advertise services through social media, calling potential victims, or even doing door-to-door sales. Also, word of mouth by victims can lead to additional fraud: “Hey, I got this great deal on cable from this guy! Saved me lots of money. You should check them out. Here’s the phone number.”

## The Business Context



**Figure 1 - Fraud in the Business Context**

It's common to believe some of these frauds are victimless crimes. That is simply not the case. Service provider losses can be extensive and include theft of service, theft of equipment, loss of funds (through, for example, refund fraud), etc. Moreover, detecting, investigating, mitigating, and preventing fraud is quite expensive for operators. There are obvious victims, of course: victims of identity theft, stolen credit card information, cloned phones or cable modems, and more. One of the important aspects to keep in mind about frauds in this space is that they are usually done en masse. A given fraudster is likely to conduct hundreds or thousands of these frauds at a time, possibly against multiple service operators.

## 4. Fraud vs Cyber Security

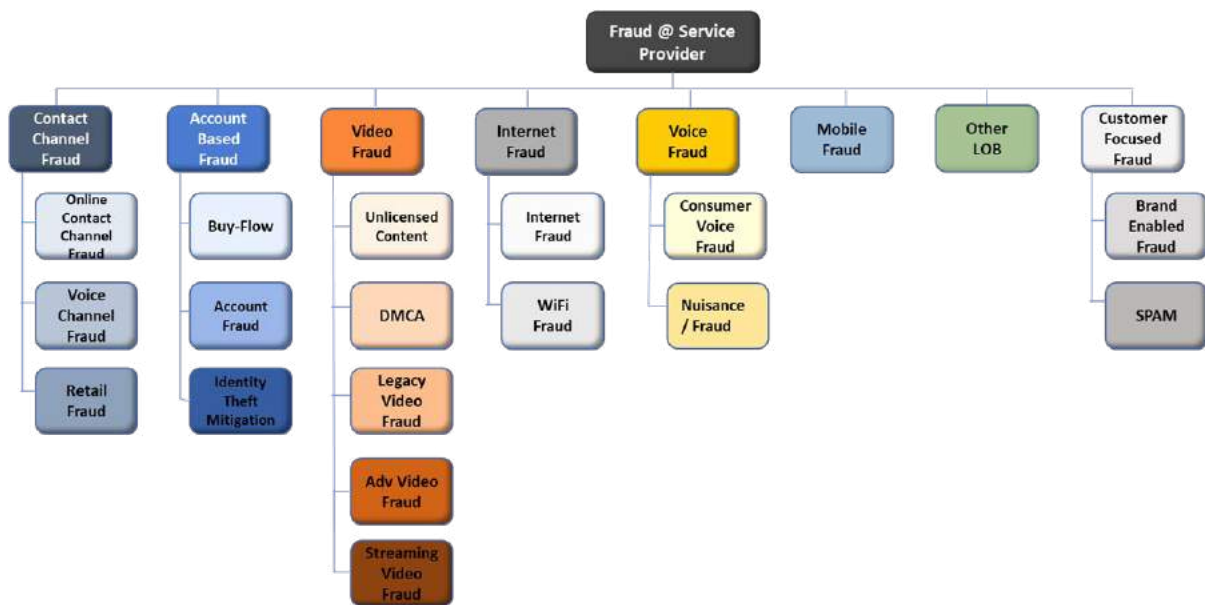
There are many types of fraud, some very simple, some very complex. The Association of Certified Fraud Examiners (ACFE) defines fraud rather broadly: "In the broadest sense, fraud can encompass any crime for gain that uses deception as its principle modus operandus." Fraud involves, at some level, misrepresentation of fact in the course of conducting crime. This is an important distinction, in comparison to cybercrime. Cyber-attacks may, or may not, involve fraud. Fraud against an operator or its subscribers may or may not benefit from illegally obtained cyber or private information (though they often do).

One of the challenges that fraud and cyber teams face is the overlap between the work of the fraud professional and the work of the cyber security professional. In the hands of a technology savvy fraudster, the fraudster will often start their fraud attack as a cyber security attack. The cyber security attack, or a series of cyber security events, will end in a fraud event. For an operator whose products are primarily technology devices and services delivered via a technology base, the lines between what is a fraud event and a cyber security attack can be blurred. Given this context, there needs to be much greater collaboration between the cyber security teams and the fraud teams.

## 5. A Fraud Framework

An important distinction from the ACFE is that there are three primary types of fraud. *Internal fraud* is conducted typically by insiders and benefiting by the access those employees, partners, or contractors have to proprietary and unique resources. This is also known as *occupational fraud*. Another type is *external fraud* which is conducted by outsiders. Another large-scale distinction, *individual fraud*, is that fraud may be targeted against specific individuals. Operators can be impacted by individual fraud when their employees may be targets of individual fraud as part of attempts to defraud the operator. Fraudsters may also use an operator's services, reputation, or information illicitly to conduct individual fraud against subscribers. All three of these types of fraud are important to network operators.

In the business of a system operator, there are many products, and the technology platforms for service and device delivery are very complex. The "Fraud Framework" illustrated in Figure 2 provides a taxonomy for examining the customer contact points through which fraud will flow, as well as specific product areas, each with its own fraud challenges.



**Figure 2 - A Fraud Framework**

In collaboration between service providers, several fraud categories have been identified. Here is a list of only a few:

- **Contact Channel Fraud** – This is fraud that has as an entry point to the contact channels the service provider has created for its customers. The contact channel type will determine the class of fraud that is attempted through a contact channel. For the online contact channel, the attacks include credential stuffing, brute force attacks, and other cyber-based incidents. The voice contact channel will experience fraud attempts using social engineering as the primary tools. The retail contact channel will include social engineering, identity theft, synthetic identity, and cyber-attacks as the tools of fraud.
- **Account or Subscriber Fraud** - As described in the introduction, compromised private information, such as a Social Security number and payment form, are used to establish an account with a service provider. The payment method is fraudulent and is used to bypass the deposit

requirements. The fraudster resells the account with an end-user who pays for access. The fraud in this case is usually a representation, by the fraudster, of somebody else. Sometimes, there can be multiple frauds – the fraudster represents to the legitimate cable operator that they are a given party using stolen personal data, and represents to the end customer that they are an authorized agent or even employee of the cable operator.

- **Video Fraud** – Video fraud can manifest itself in many forms, including unlicensed content, notices from the Digital Millennium Copyright Act (DMCA) that the operator needs to process, legacy and advanced set-tops being used for illegitimate video service delivery, and finally streaming fraud enabled by credential sharing or credential fraud. There are many ways to monetize stolen content -- probably as many as there are to steal the content in the first place. The fraud here is that people distribute property and take payment for which they don't have the rights. Like voice fraud, there are several industry groups seeking to address video piracy and make the associated frauds simply unprofitable.
- **Internet Fraud** – Internet fraud includes the cloning or hacking of cable modems, as well as the illegitimate use of WiFi hotspots provided by the service provider.
- **Voice Fraud and Nuisance Voice** - There are many frauds conducted using fixed and mobile telephone networks. These range from toll fraud to spam to phishing and more. There are already several government agencies in many countries and group forums where operators, law enforcement, and regulators collaborate to address voice fraud. Voice fraud is obvious when someone calls a victim and misrepresents themselves. However, spoofing an originating phone number belonging to a known brand that is not owned by the fraudster to connect to a consumer is an effective way of committing fraud against unsuspecting consumers.
- **Mobile or Cell Phone Fraud** - This includes subscriber fraud, as described previously; it also includes cell phone cloning fraud. A common example today is SIM hijacking or SIM swapping, where a subscriber identity module (SIM) uses credentials from a legitimate subscriber's mobile phone's SIM card to activate one or more other cell phones. Because of the need for mobile service to support roaming, mobile networks have unique vulnerabilities to fraud that may impact multiple operators.
- **Customer-Focused Fraud** – in this class of fraud, the brand of the operator is used to gain confidence with the consumer either via voice contact or email. The confidence is used to defraud the consumer.

## 6. Details Relevant for Explaining Specific Frauds

Fraud is complicated. Many frauds are described by the Association of Certified Fraud Examiners and other groups. Two particularly good resources for understanding frauds are the Open Risk Manual [4] and the “Framework for a Taxonomy of Fraud” [5] developed by Stanford Center on Longevity, in collaboration with the Financial Industry Regulatory Authority (FINRA). Yet, cable operators experience frauds that are not even described by these resources. Our adversaries are very creative at finding ways to extract funds from our companies and our customers.

Can frauds be categorized and described in detail? What are the details that would be useful to cyber security and fraud professionals in describing types of fraud or related cybercrimes?

Like the frauds themselves, this gets complicated. However, an ontology and a taxonomy can be developed to describe and communicate about frauds. These details start with just basic descriptions and understanding the goal of the fraudsters. It's also useful at this level of detail to specifically identify the fraud that is being performed – exactly how is the criminal misleading their target, misrepresenting themselves, or otherwise specifically breaking the law. Who are the targets of the frauds and what type of fraudsters execute the fraud? (Yes, there are different skill sets and, as a consequence, different types of

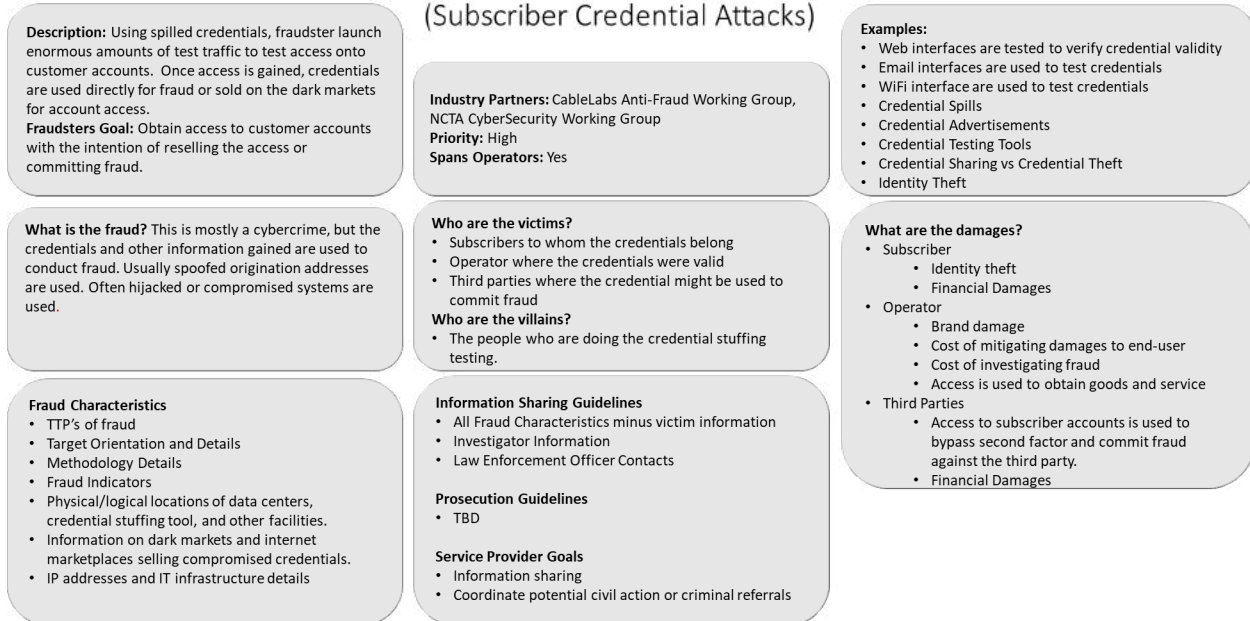
fraudsters.) It is useful to specify the damages to customers, the operator, or other parties (such as the government or vendors). Then there are details of the fraud methods themselves. These may include:

- Tactics, techniques, and procedures (TTPs) used to conduct fraud (TTP is also a cybersecurity term)
- Target Orientation and Details
- Methodology Details
- Fraud Indicators
- Physical/logical locations of data centers and other facilities
- Hacking tools used to execute the fraud (credential stuffing tools or markets, phishing frameworks, social engineering tools, etc.)
- Information on dark markets and Internet marketplaces selling compromised credentials or other enabling information for the fraud
- IP addresses and other IT infrastructure details

Then there are information elements useful to the examiners and investigators researching the fraud: Priority of the investigation, interested industry partners (perhaps other operators also being attacked), information sharing guidelines (what can or cannot be shared outside a company, for example). There are details related to enforcement actions such as prosecution guidelines, law enforcement or regulator contacts, and identities of investigators themselves. Finally, it is useful to understand the statistics and trends of the given fraud. This may include damages per event, cost to mitigate per event, estimated number of events annually, known cases, and the scale of the fraudster ecosystem perpetrating the fraud. Collected over time, these details can be used to identify trends.

This information can be collated and then presented in a very concise manner to quickly describe a type of fraud. For example, consider the description in Figure 3, below, of subscriber credential attacks as part of the “Online Contact Channel” category. In one page, specific information is presented that specifically conveys what a subscriber credential attack looks like. This is, of course, simply illustrative and lacks details of an actual fraud.

## Online Contact Channel Fraud (Subscriber Credential Attacks)



**Figure 3 - A Taxonomy Example of the Online Contact Channel Fraud**

## 7. What Use is a Taxonomy?

A taxonomy, defined loosely here as a classification of fraudulent activities, can be very useful. It provides a way to specifically describe a fraud, which can be useful for deciding how to respond, and possibly even to support workflow automation. The “Framework for a Taxonomy of Fraud” [5] provides an excellent example of how this might be done by actually codifying fraud. Their system formalizes the ideas above and uses codes to classify the fraud and adds tags for describing the incident, victim, and perpetrator. An example they show in their report is “1.7.1. AD:IE. PS:M. MT:PC. FV. EF. MP.” This string of characters is translated as:

1.7.1-> Classification Number

1->Individual Financial Fraud

7->Relationship & Trust Fraud

1-> Romance Scam

AD:IE -> Method of Advertising Fraud(AD): Internet, email(IE)

PS:M -> Purchase Setting(PS):Mail(M)

MT:PC -> Method of Money Transfer(MT): Personal Check(PC),

VT:FV, EF -> Victim Tags(VT):Female Victim(MV), Elder Victim 65+((EV)

PT:MP -> Perpetrator Tags(PT):Male Perp(MP)

While it may not always be beneficial to develop or use such a proforma format, the notion of specifically describing fraud is clearly useful for ensuring accurate communications between professionals investigating and prosecuting frauds. A taxonomy allows very clear discussion of exactly the nature and method of fraud. This can help operators solicit input from other operators to see if a given actor is conducting similar fraud in other markets, for example.

We've also found that a taxonomy can help to provide an organizational structure for how operators respond to fraud. Once they've codified what frauds they experience, they can organize accordingly. Moreover, the taxonomy can be useful for ensuring resource allocation and that workflows are optimized to deal with specific frauds.

## 8. Conclusion

Fraud is a challenge to all operators and will continue to be. It has a large impact to services providers who are at risk for both financial fraud and technology enabled fraud. A fraud taxonomy provides a way to classify a complex fraud system and improves the industry's ability to communicate, discuss, and improve the service provider's fraud detection and prevention capabilities. While an individual fraud attempt aimed at a service provider is targeted to that operator, the tools, techniques, and procedures (TTP) are common and shared between the fraudsters. The fraudsters share TTP's and so should we as an industry. The authors are seeking collaborators to take this work to greater levels of detail.

## Abbreviations

|       |                                          |
|-------|------------------------------------------|
| ACFE  | Association of Certified Fraud Examiners |
| DRM   | Digital Rights Management                |
| FINRA | Financial Industry Regulatory Authority  |
| ID    | Identification                           |
| IVR   | Interactive Voice Response               |
| LTE   | Long Term Evolution                      |
| SIM   | Subscriber Identification Module         |
| TTP   | Tactics, Techniques and Procedures       |

## Bibliography & References

- [1] "The Top 11 Fraud Types Here in the US and Their Losses", online: <https://frankonfraud.com/fraud-trends/the-top-fraud-losses-for-2019-by-fraud-type/>
- [2] OWASP Credential Stuffing, online: [https://owasp.org/www-community/attacks/Credential\\_stuffing](https://owasp.org/www-community/attacks/Credential_stuffing)
- [3] "The Economy of Credential Stuffing Attacks", online: <https://www.recordedfuture.com/credential-stuffing-attacks/>
- [4] Open Risk Manual, online: [https://www.openriskmanual.org/wiki/Main\\_Page](https://www.openriskmanual.org/wiki/Main_Page)
- [5] "Framework for a Taxonomy of Fraud", Stanford Center on Longevity and FINRA, 2015, online: <http://longevity.stanford.edu/2015/07/30/framework-for-a-taxonomy-of-fraud/>



# **Fraud Prevention and Privacy Law**

## **Emerging Conflicts Between Privacy Law and Fraud Prevention**

A Technical Paper prepared for SCTE•ISBE by

**Will Bracker**

Corporate Counsel, Privacy  
Cox Communications  
Atlanta, GA  
Will.Bracker@cox.com

**Steve Goeringer**

Distinguished Technologist  
CableLabs  
Louisville, CO  
s.goeringer@cablelabs.com

**Simon Krauss**

Deputy General Counsel  
CableLabs  
Louisville, CO  
s.krauss@cablelabs.com

# Table of Contents

| <b>Title</b>                                                         | <b>Page Number</b> |
|----------------------------------------------------------------------|--------------------|
| 1. Introduction.....                                                 | 3                  |
| 2. Comprehensive Privacy Laws .....                                  | 3                  |
| 3. Comprehensive Privacy Law vs. Fraud Detection and Prevention..... | 5                  |
| 3.1. Right to know.....                                              | 6                  |
| 3.2. Right to be forgotten .....                                     | 6                  |
| 4. Fraud Prevention and Privacy Law in a Pandemic.....               | 7                  |
| 5. Conclusion .....                                                  | 8                  |
| Abbreviations.....                                                   | 8                  |
| Bibliography & References .....                                      | 8                  |

## List of Figures

| <b>Title</b>                                                               | <b>Page Number</b> |
|----------------------------------------------------------------------------|--------------------|
| Figure 1: Status of US state level comprehensive privacy legislation ..... | 4                  |

# 1. Introduction

Laws and regulations protecting privacy are not new. In Western civilization, case law on privacy extends back to the early 1400's, when the law prohibited eavesdropping. But what is new is the emergence of comprehensive privacy laws and their supporting regulations. These statutory schemes create broad rights for citizens and impose significant obligations on businesses with respect to collection, use and protection of personal information. Steep non-compliance penalties and short implementation timelines require businesses to build robust compliance programs.

In a global economy, building these new compliance functions is no easy task: each new comprehensive privacy law is different from the last, creating a challenging environment for multi-state or multi-national enterprises. New rights and obligations also have the potential to provide new lines of attack for fraudsters and can limit the ability to detect and prevent fraud. Finally, the global pandemic brings a heightened challenge and creates even more opportunities for bad actors to compromise privacy and evade consequences.

This paper examines two specific privacy requirements in light of an operator's need to conduct fraud detection, mitigation, investigation, and prevention. Considerations include anti-fraud information collection, sharing, and action. These areas, indeed anti-fraud operations in general, are often overlooked as risk and legal departments draft compliance program guidelines.

## 2. Comprehensive Privacy Laws

On May 25, 2018 the General Data Protection Regulation (GDPR) came into effect in the European Union, ushering in a new era for privacy protection: The Age of the Comprehensive Privacy Law. Rather than addressing privacy and data protection in an individual business sector or activity, a comprehensive privacy law grants citizens global control over the collection and use of their personal information. Broadly speaking, these rights allow citizens to access their data, understand how it is used and who it is being shared with, correct errors, restrict use, and require deletion.<sup>1</sup> New business obligations are also part of the landscape, generally requiring greater transparency as to data held by the business, limitations on processing and use, and enhanced duties on those who would use the personal data of their citizens.<sup>2</sup>

These laws require new regulations and new regulators, which in turn requires that businesses create new compliance programs to ensure that they do not fall afoul of another commonality of these laws: substantial penalties for non-compliance.<sup>3</sup>

In the United States, comprehensive privacy law at the state level is moving quickly. California's Consumer Privacy Protection Act of 2018 (CCPA) was the first such law to see the light of day – but more are coming. According to the International Association of Privacy Professionals (IAPP), sixteen

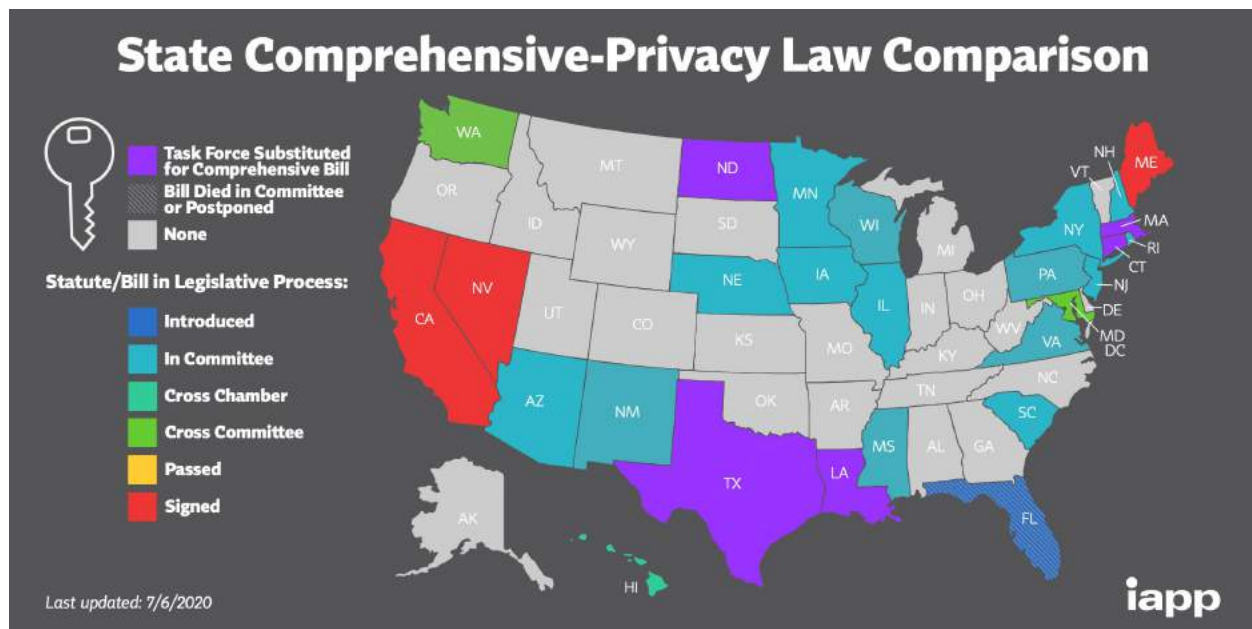
---

<sup>1</sup> The International Association of Privacy Professionals (IAPP) has identified commonalities that apply to these privacy laws: Access to collected data, access to shared data, right to correct data, delete data, restrict the use of personal data, data portability, the right to opt out of use, prohibit automated decision making, and whether or not a consumer can sue for violations of the law, also referred to as a private right of action.

<sup>2</sup> These obligations as identified by IAPP include: Age based opt-in, notice and transparency requirements, data breach notifications, risk assessments, prohibitions on discrimination, purpose limitations, processing limitations, and heightened or fiduciary duties for storage and use of personal data.

<sup>3</sup> General Data Protection Regulation, Chapter VII, Article 83; California Consumer Privacy Act of 2018 1798.150, .155

states have a comprehensive privacy bill at some stage of their legislative process. In addition to California (which is likely to revise the 2018 CCPA Act, which just became effective, via the California Privacy Rights Act of 2020), Maine and Nevada have already passed legislation. Other states with legislation in progress are Arizona, Connecticut, Hawaii, Illinois, Iowa, Louisiana, Maryland, Massachusetts, Minnesota, Nebraska, New Hampshire, New Jersey, New York, North Dakota, South Carolina, and Texas. Many of these states have more than one bill under consideration. Legislation was raised in eleven other states. This is summarized in Figure 1 (used with permission of IAPP). See the IAPP website for current details (<https://iapp.org/resources/article/state-comparison-table/>).



**Figure 1 - Status of US state level comprehensive privacy legislation**

Simply put, merely keeping up with all of the state legislative activity is a daunting task.

One of the challenges operators face is that different jurisdictions (*e.g.*, Europe, Canada, different states, etc.) have enacted and are in the process of drafting and/or enacting different laws, so there will be multiple frameworks to apply depending on the context of the private information being protected and the regions in which it applies or exists. For example, consider a seemingly simple concept like the definition of personal information. The GDPR defines personal data as:

any information relating to an identified or identifiable natural person (‘data subject’);

PIPEDA is equally brief in its definition of personal information:

information about an identifiable individual.

CCPA is considerably more verbose:

“Personal information” means information that identifies, relates to, describes, is capable of being associated with, or could reasonably be linked, directly or indirectly, with a particular consumer or household.

in fact, the full text of the CCPA definition extends to some 238 words, reproduced in the footnote below.<sup>4</sup>

The regulatory process that follows the enactment of the law adds to the complexity of the compliance landscape. The rulemaking process for the CCPA recently came to a close, six months after the underlying law came into effect.

In California, the cost of initial compliance with the law was estimated at a staggering \$55 billion. Depending on the number of firms ultimately deemed to be covered by the CCPA, the ongoing annual compliance is estimated to be between \$466 million and \$16.4 billion.

In addition, once laws are enacted, court interpretations of those laws will doubtless vary. This paper provides examples of the challenges of compliance with the variety of privacy laws found in the CCPA (which may change again in November), the GDPR (which, while focused on data protection, does address privacy), and Canada's Personal Information and Electronic Documents Act (PIPEDA).

Complexities aside, the intent of a comprehensive privacy law is to allow a citizen greater visibility and control over the use of their personal data and to require businesses to focus attention and resources on the protection of that personal information. However, as written, these laws may actually make that task more difficult.

### 3. Comprehensive Privacy Law vs. Fraud Detection and Prevention

Fraud (n): *deceit, trickery*. Wrongful or criminal deception intended to result in financial or personal gain.<sup>5</sup>

The data ecosystem of privacy law between company and citizen customer has a third actor lurking in the shadows: fraudsters. Such individuals, criminal rings, and nation-state actors all use the same basic

---

<sup>4</sup> (o) (1) "Personal information" means information that identifies, relates to, describes, is capable of being associated with, or could reasonably be linked, directly or indirectly, with a particular consumer or household. Personal information includes, but is not limited to, the following:

(A) Identifiers such as a real name, alias, postal address, unique personal identifier, online identifier Internet Protocol address, email address, account name, social security number, driver's license number, passport number, or other similar identifiers.

(B) Any categories of personal information described in subdivision (e) of Section 1798.80.

(C) Characteristics of protected classifications under California or federal law.

(D) Commercial information, including records of personal property, products or services purchased, obtained, or considered, or other purchasing or consuming histories or tendencies.

(E) Biometric information.

(F) Internet or other electronic network activity information, including, but not limited to, browsing history, search history, and information regarding a consumer's interaction with an Internet Web site, application, or advertisement.

(G) Geolocation data.

(H) Audio, electronic, visual, thermal, olfactory, or similar information.

(I) Professional or employment-related information.

(J) Education information, defined as information that is not publicly available personally identifiable information as defined in the Family Educational Rights and Privacy Act (20 U.S.C. section 1232g, 34 C.F.R. Part 99).

(K) Inferences drawn from any of the information identified in this subdivision to create a profile about a consumer reflecting the consumer's preferences, characteristics, psychological trends, preferences, predispositions, behavior, attitudes, intelligence, abilities, and aptitudes.

<sup>5</sup> <https://www.lexico.com/en/definition/fraud>

approach to access the value contained within that ecosystem: they lie about their identity. Using a combination of personal data gathered from public and stolen private sources as well as purely synthetic identities, fraudsters attempt to deceive one side of the transaction into believing that they are the other expected actor, and then extract the value in the transaction for themselves.

Preventing such attempted deception is at once simple and complex: the key lies in having sufficient personal information about the other actor in the transaction to distinguish a fraudster from a legitimate actor. Techniques for accomplishing this range from a simple picture ID check by a retail clerk to sophisticated big data solutions that evaluate the risk of a transaction based on hundreds or thousands of data points processed by sophisticated algorithms. Simple or complex, the core component of every prevention technique is the use of personal information to disprove the fraudster's impersonation of the legitimate actor.

Comprehensive privacy laws can be used as both sword and shield to a fraudster. Offensively, the right to know can be used to gain information about an individual to make impersonation harder to detect. Defensively, fraudsters can use the transparency, opt-out, and deletion rights of the new laws to hinder or evade detection and to determine what data is being used to detect them.

### **3.1. Right to know**

Citizen rights under comprehensive privacy laws start at a common point: in order to exercise the other rights granted, it is important for the citizen to know what personal data a given company is storing about them. Therefore, the first right granted under a CPL is the right for a citizen to require a company to provide a record of that data to the citizen. This represents a brand-new opportunity for fraudsters to accumulate the information necessary to impersonate their targets.

How this is done was documented by James Pavur and Casey Knerr at Blackhat 2019. In their seminal white paper, *GDPArrrrr: Using Privacy Laws to Steal Identities*, they show the weaponization of privacy tools and laws, using GDPR "right to know" requests on 75 companies, and providing only publicly available information about the subject as authentication. The GDPR right to know is one of the basic data subject rights; willful failure to comply with this requirement can subject a company to a penalty up to the greater of 20 million euros or 4% of the firm's annual worldwide gross revenue. Given the severity of the penalty for non-compliance, it is perhaps unsurprising that some companies erred on the side of compliance and divulged at least a subset of the requested personal information.

James and Casey then used the data from the responses to craft 75 additional right to know requests. Enriched with the data from their initial harvest, these requests were much more productive. They received 10 digits of credit card numbers, expiration date, login credentials for websites, educational test scores, complete hotel records, dating profiles, purchase histories, and much more.

### **3.2. Right to be forgotten**

Fool me once, shame on you. Fool me twice, shame on me... but what if I had to forget that you fooled me?

Some of the citizens' rights and business obligations created by new comprehensive privacy laws are being misused by bad actors to both further their fraud schemes and to escape or evade fraud detection. Fraudsters can leverage new consumer rights that commonly include access to collected data, access to shared data, correct data, deleting data, opting-out of use, and prohibited automated decision making. Business obligations can also be abused including notification and transparency requirements, purpose limitation, and processing limitations.

For example, under GDPR, bad actors use the right to know in order to gather personal information of their targets. They also have used the right to be forgotten to conceal their identities and reduce the effectiveness of notice, public safety, and fraud prevention tools.

Part of the challenge is that consumers don't have a specific, immutable identifier on which operators can tie personal information. Thus, it's very hard implement consumer interfaces to provide privacy management features that are secure – that is, hard to misuse.

In addition, privacy laws limit fraud investigations by limiting both what information is collected and how an operator may use it. For example, GDPR contains a limited exception to share information without consent when “processing is necessary for the purposes of legitimate interests pursued by the controller or by a third party”. Preventing fraud has been defined as a legitimate interest. However, the sharing exception language was not carried over into the subscriber's right to have their information deleted.

The CCPA also includes several exceptions for the collection, retention, and disclosure of personal information that intersect with anti-fraud efforts. These are:

- To detect data security incidents, or protect against fraudulent or illegal activity;
- To comply with laws;
- To comply with government investigations;
- To cooperate with law enforcement; and
- To exercise the defense of legal claims, including evidentiary privilege.

Moreover, there are additional areas that must be considered that are not well defined within current laws. Privacy is transitive. So, as fraud investigations are conducted, while information that is shared with law enforcement and other governmental agencies may be in compliance with governmental investigations, does information shared with private entities that may benefit from receiving the information in relation to their own fraud investigations (such as other operators) give rise to privacy law violations? Other laws may also govern the disclosure and use of private information, such as the U.S. Fair Credit Reporting Act, which may govern how an operator may use information about prior fraudulent activity to deny a person service.

An awkward question is whether fraudsters themselves enjoy privacy protection. A clear example of contention, for example, is the case where fraudsters are also customers or even business partners (such as channel sales companies). However, many cases are not so clear. Is a fraudulent customer's personal data (which may not even be legitimate) protected? Is personal information from cyber-attacks protected?

## **4. Fraud Prevention and Privacy Law in a Pandemic**

Finally, the COVID-19 pandemic has made the task of fraud prevention more complex while simultaneously raising the payoff for fraudsters.

Distanced communication makes the process of investigation more difficult. Interviews are being conducted via video or teleconference, depriving investigators of the ability to observe their subjects. In the pre-COVID world, many exception processes to web- or phone- based provisioning or identity flows would direct the individual to go to a physical location and present themselves with one or more forms of government issued identification for visual comparison by a human being. But in the present circumstances, consumers are unable to go to a physical service location and interact with an in-person

representative of an agency or a business. All parties in a transaction are instead left reliant on telephone or digital verification methods to measure the risk of a transaction and decide whether or not to proceed.

This weakening of in-person verification processes has not gone unnoticed by fraudfeasors. The FTC reports that, as of June 28<sup>th</sup>, U.S. citizens have reported losses more than \$108 million to COVID-19 related scams and fraud schemes. Identity theft of both individual and business entities is on the rise as bad actors take aim at enriched unemployment benefits and small business support programs.

## 5. Conclusion

This article provides just a short survey of the issues that intersect privacy and fraud. It does not purport to provide any prescriptive approaches or formulas for how to address those issues. That is because it cannot do so—such solutions are simply too complex. Moreover, compliance is jurisdictionally specific. The requirements and obligations that drive compliance for one operator will not fully apply to another operator operating in states, let alone countries. Nevertheless, it is important that compliance teams consider fraud detection, investigation, mitigation, and prevention efforts as they develop and evolve guidelines and procedures for their companies.

## Abbreviations

|        |                                                        |
|--------|--------------------------------------------------------|
| CCPA   | The California Consumer Privacy Act                    |
| GDPR   | The General Data Protection Regulation                 |
| PIPEDA | The Personal Information and Electronic Documents Act  |
| IAPP   | The International Association of Privacy Professionals |

## Bibliography & References

California Consumer Privacy Act, State of California Department of Justice, online, <https://oag.ca.gov/privacy/ccpa>

Data protection in the EU, European Commission, online, [https://ec.europa.eu/info/law/law-topic/data-protection/data-protection-eu\\_en](https://ec.europa.eu/info/law/law-topic/data-protection/data-protection-eu_en)

The Personal Information Protection and Electronic Documents Act (PIPEDA), Office of the Privacy Commissioner of Canada, online, <https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/>

US State Comprehensive Privacy Law Comparison, IAPP, online, downloaded July 20, 2020. <https://iapp.org/resources/article/state-comparison-table/>

“GDPArrrrr: Using Privacy Laws to Steal Identities”, James Pavur and Casey Knerr, Blackhat USA 2019 White Paper.



# Approaches to Latency Management: Combining Hop-by-Hop and End-to-End Networking

A Technical Paper prepared for SCTE•ISBE by

**Sebnem Ozer, Ph.D.**

Senior Principal Architect  
Comcast  
1800 Arch St., Philadelphia, PA 19103  
2152868890  
Sebnem\_Ozer@comcast.com

**Carl Klatsky**

Senior Principal Engineer, Product Development  
Comcast  
1800 Arch St., Philadelphia, PA 19103  
215-286-8256  
Carl\_Klatsky@comcast.com

**Dan Rice**

VP  
Comcast  
1401 Wynkoop St Ste 300, Denver, CO 80202  
720-512-3730  
Daniel\_Rice4@comcast.com

**John Chrostowski**

Executive Director  
Comcast  
1800 Arch St., Philadelphia, PA 19103  
267-260-3695  
John\_Chrostowski@comcast.com

# Table of Contents

| Title                                                                      | Page Number |
|----------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                       | 3           |
| 2. Latency Measurement.....                                                | 3           |
| 3. Latency Performance in Current and Emerging Network Architectures ..... | 10          |
| 4. End-to-end Support For Low Latency Services .....                       | 17          |
| 5. Conclusion: Final Thoughts on Latency Management.....                   | 20          |
| Abbreviations .....                                                        | 22          |
| Bibliography & References.....                                             | 23          |
| Acknowledgments .....                                                      | 23          |

## List of Figures

| Title                                                                                                                                                                                                                                                                                                                                | Page Number |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| Figure 1 – Examples of Latency Measurement Methods .....                                                                                                                                                                                                                                                                             | 5           |
| Figure 2 – TCP connection based latency measurement. Top: Internet RTT; Bottom: Client RTT .....                                                                                                                                                                                                                                     | 8           |
| Figure 3 – Latency Under Load Measurement .....                                                                                                                                                                                                                                                                                      | 9           |
| Figure 4 – Top Left: Max DS and US LUL with suboptimal DS AQM settings; Bottom Left: Mean DS and US LUL with suboptimal DS AQM settings; Top Right: Max DS and US LUL with optimized DS AQM settings; Bottom Right: Mean DS and US LUL with optimized DS AQM settings .....                                                          | 12          |
| Figure 5 – Top: Max US LU; Bottom: Mean US LUL for different CM models and speed tiers.....                                                                                                                                                                                                                                          | 13          |
| Figure 6 – Left: dslreports.com results with suboptimal DS AQM settings and D3.1 CM with BC with 250ms default target latency; Right: Optimized DS AQM settings and D3.1 CM with US AQM with 10ms default target latency .....                                                                                                       | 14          |
| Figure 7 –Throughput (Mbps) values over time for each flow per CM HSD buffer size. Each flow has different e2e RTT. Green flow's server is the closest to the subscriber's home while purple flow's server is the farthest away. Top: 10ms target buffer size; Middle: 30ms target buffer size; Bottom: 50ms target buffer size..... | 15          |
| Figure 8 – End-to-end LL services support.....                                                                                                                                                                                                                                                                                       | 17          |
| Figure 9 – Low Latency Services Monitoring and Management for CCAP Systems.....                                                                                                                                                                                                                                                      | 19          |
| Figure 10 – Low Latency Services Monitoring and Management for Distributed Systems .....                                                                                                                                                                                                                                             | 19          |
| Figure 11 – Low Latency Services Monitoring and Management Integrated within Data-driven and Knowledge-defined Architectures .....                                                                                                                                                                                                   | 20          |

## List of Tables

| Title                                                            | Page Number |
|------------------------------------------------------------------|-------------|
| Table 1 – Properties of Latency Measurement Methods .....        | 6           |
| Table 2 – DOCSIS Networks Latency: RTT Between CM and CMTS ..... | 10          |
| Table 3 – Architecture changes to support LL services .....      | 18          |

## 1. Introduction

MSO networks that deliver services between the source and destination ends consist of multiple hops, i.e. different network segments. Customer experience is shaped by many Quality of Service (QoS) features that may be defined per hop-by-hop and end-to-end views. Traditionally, speed performance has been addressed as the main contribution to Quality of Experience in an MSO network. However, latency and jitter have a significant impact on many current and emerging services. MSOs have started modeling, monitoring and managing a more complete performance concept, including speed, latency, jitter, packet loss, security and reliability. The importance of this approach has been reiterated in lockstep with a significant increase in traffic volumes, across a very different mix of residential services, and because of a sudden pandemic, as MSOs who adopted this approach were successful supporting additional traffic volumes. This paper will describe current work on latency measurement, optimization and management platforms with hop-by-hop and end-to-end features. We will present current achievements and results for lower latency systems in the cable industry, and ongoing work on new optimization techniques and big data analytics. Architectural examples will be provided for a data-driven and knowledge-based converged access network with low latency service assurance and agility. Finally, we will discuss roadmap items MSOs may adopt to provide low latency services within their 10G initiative.

## 2. Latency Measurement

Subscribers' Quality of Experience is a subjective concept affected by the Quality of Service along the path between the service endpoints. QoS metrics are defined from the system's perspective and can be measured and managed. QoE requires a multi-disciplinary approach to assess the user's perspective that may be affected by factors unique to the user and the user's interactions with the service. Mapping QoS to QoE is still evolving as new services and applications emerge. Performance metrics such as throughput (speed) have been regularly measured by MSOs, but speed is only one of the performance indicators. Different services and applications require different levels of speed, latency, jitter and packet loss. For example, depending on the online gaming type and platform, the impact of latency, jitter and packet loss on the gamer's experience and his/her lag perception may be different [1]. Customer experience with gaming, videoconferencing, VR/AR and many commercial services can only be assessed if all the corresponding QoS metrics -- and especially latency and jitter -- are also well modeled, monitored and managed. Fairness, availability, reliability and security are other factors affecting customer satisfaction.

MSOs need to assess and optimize the latency/jitter performance in their networks to improve their subscribers' experiences, and to ensure a strong competitive position, as intended by the 10G initiative. Upstream and downstream usage changed, due to the Covid-19 pandemic, and the steady increase of low latency services in residential networks is increasing the priority of latency management systems.

Figure 1 shows latency measurement (LM) methods widely used by operators. The LM-1 method is a passive measurement method described in more detail later in this section. It measures the round-trip time (RTT) of Transmission Control Protocol (TCP) handshakes to assess access and internet hops, as well as end-to-end latency values (including home, core networks and internet.) It has the advantage of using actual production traffic with no impact on the network. Due to TCP's ACK suppression and expediting implemented in CMs, it may not be directly mapped to UDP type latency. This latency is affected by network conditions and utilization levels (e.g. media access delay due to utilization in DOCSIS bonding groups, or Wi-Fi channel and queueing delay due to HSD service consumption.)

The LM-2 method is an active measurement for upstream (US) and downstream (DS) traffic latency under load (LUL). The test starts with TCP-based connection establishment only if the gateway utilization is low because of potential customer-facing performance impacts. This TCP-based connection latency within LM-2 (conn RTT in Figure 1 and Figure 3) reflects the performance in idle conditions. After the successful connection, LUL in both directions are measured sequentially. US LUL is round trip latency under US load while DS LUL is round trip latency under DS load and they reflect the queuing delay ranges in the CM and CMTS respectively. The LUL test may use iperf TCP as the load to utilize the bandwidth up to speed tier rates in DS and US directions. UDP pings (e.g. Netperf UDP\_RR or iRTT) are transmitted concurrently. All the additional test traffic is excluded from any billing. UDP ping latency, jitter and packet loss metrics are collected along with TCP load throughput. Any bufferbloating issue can be detected and the performance of buffer control and active queue management techniques can be optimized with this test method. It can be measured at different endpoints to assess home Wi-Fi and access network performance.

The LM-3 method comprises well known ping measurements (both ICMP and DOCSIS pings). ICMP pings can be used when data and control planes overlap, but are not suitable for SDN networks where data and control planes are not the same. ICMP pings may be routed differently than TCP/UDP packets and may be processed with lesser priority. This latency is also affected by utilization levels similar to LM-1 method, however these are periodic synthetic ping data.

An overview of the properties of these measurement methods is provided in Table 1. LM-1 and LM-2 methods are newer techniques that MSOs started to integrate into their performance measurement platforms. More detailed information on LM-1 and LM-2 methods with test and measurement points that can be integrated are described below.




Figure 1 – Examples of Latency Measurement Methods

**Table 1 – Properties of Latency Measurement Methods**

| Test/<br>Monitoring                                                | Traffic End<br>Point 1       | Traffic End<br>Point 2  | Monitoring /<br>Test Points                     | Traffic                                            | Other<br>metrics                                                                       | Comments                                                                                                                                                                                                                                                                                   |
|--------------------------------------------------------------------|------------------------------|-------------------------|-------------------------------------------------|----------------------------------------------------|----------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>LM-1</b><br><br><b>Passive<br/>Measurements</b>                 | CPE                          | Internet                | SM-VNF                                          | TCP<br>traffic                                     | <ul style="list-style-type: none"> <li>•TCP session info and retx count</li> </ul>     | <ul style="list-style-type: none"> <li>•No extra traffic &amp; End-to-end approach</li> <li>•TCP traffic latency, no UDP latency</li> <li>•US latency is impacted by TCP suppression and expediting</li> <li>•Capability to measure per SF performance and home HSD utilization</li> </ul> |
| <b>LM-2</b><br><br><b>Active<br/>Measurements<br/>Under Load</b>   | CPE or gateway erouter       | Netperf & Iperf Servers | Clients and server functionalities at endpoints | TCP iperf up to speed tier + Netperf UDP_RR / iRTT | <ul style="list-style-type: none"> <li>•Speed</li> <li>•Jitter, packet loss</li> </ul> | <ul style="list-style-type: none"> <li>•UDP ping traffic latency, TCP RTT</li> <li>•Should run only in idle times to not affect customer</li> <li>•Testing traffic must be excluded from billing</li> </ul>                                                                                |
| <b>LM-3</b><br><br><b>Active<br/>Measurements<br/>Without load</b> | CPE or gateway erouter or CM | CMTS or Servers         | Ping Endpoints                                  | ICMP Ping & DOCSIS MAC Ping                        | <ul style="list-style-type: none"> <li>•Jitter, packet loss</li> </ul>                 | <ul style="list-style-type: none"> <li>•ICMP/DOCSIS ping latency/jitter/packet loss</li> <li>•ICMP vs UDP/TCP diff (control plane)</li> <li>•Minimal extra traffic</li> <li>•ICMP is not directly applicable to SDN devices/vCM/vCMTS</li> </ul>                                           |

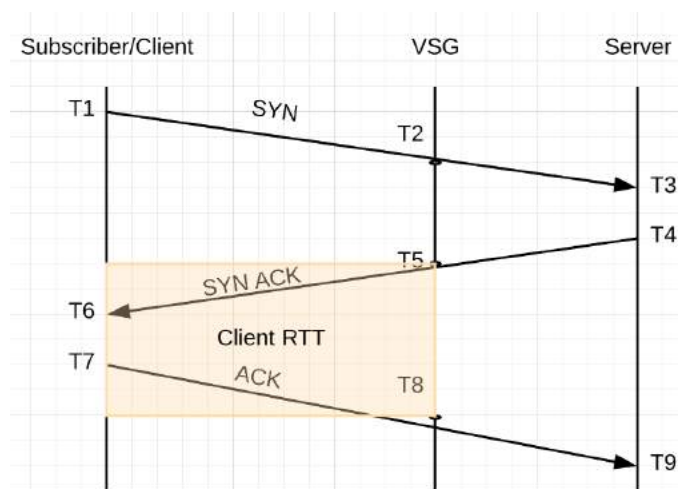

The LM-1 method has been used in Comcast Networks with Software Defined Networking (SDN) and Virtual Network Functions (VNF) components. As SDN has evolved over the last decade, the telecommunications industry has seen the opportunity that SDN presents as a means to develop new product service offerings, and specifically the means to rapidly deploy VNF in support of new products. Comcast

has realized SDN as part of its Active Core and Access platforms. Comcast has also realized SDN as part of its Virtual Services Gateway (VSG) platform. The VSG platform uses commercial off-the-shelf (COTS) compute servers, supported with appropriate network interfaces, upon which various VNFs can be instantiated in support of new products and services.

The VSG platform is deployed in line with the production data traffic and thus a variety of VNFs are envisioned, covering usage metering, monitoring, telemetry, reporting, and traffic marking. Multiple VNFs can be instantiated as supported by the platform's compute and network capability. The initial hardware iteration of the VSG platform is based on a dual-socket x86 compatible motherboard, with twenty cores per CPU socket. The VSG platform hardware also includes dual-100G interfaces, for interconnection to peer network elements. This compute & networking configuration can support the initial set of VNFs envisioned and planned.

The first VNF that we deployed on the VSG platform was a usage metering function. By offloading the existing usage metering function from the CMTS, the function becomes more versatile. That's because changes to per-user profiles can be achieved via real-time configuration, as requested through the OSS / BSS.

The second VNF that we deployed was the telemetry function introduced above. This VNF monitors the anonymized TCP connections of connected subscribers and provides a collection of TCP statistics per subscriber. Within a configurable reporting interval, the VNF is able to track things like TCP session count, TCP connection duration, total session packet count, and TCP retransmission count, per direction. Two key statistics provided by this VNF, and shown in Figure 2, are TCP SYN-to-SYNACK RTT (Server RTT) and TCP SYNACK-to-ACK RTT (Client RTT).



**Figure 2 – TCP connection based latency measurement. Top: Internet RTT; Bottom: Client RTT**

From the VSG’s position in the network, looking at the TCP SYNACK-to-ACK timing (Client RTT) specifically gives a view into the performance of the access and home portion of the network, as the other network elements in the path to the customer device are the CMTS, cable modem and Wi-Fi Access Point. The per-subscriber metrics give an anonymized view into a specific customer’s network experience. By looking at the aggregate subscriber metrics across a given service group or bonding group, we gain a view into the health of that service group or bonding group. Monitoring these metrics can serve as an early warning signal of the need for capacity expansion, even before traditional throughput and utilization metrics, thus providing a new way to plan for capacity expansion.

Providing insight into this timing serves also as a proxy for the customer experience. A part of how customers perceive the performance of the network is governed by the TCP 3-way handshake, to establish the TCP connection between the client and the server. The typical online application (Microsoft Office,

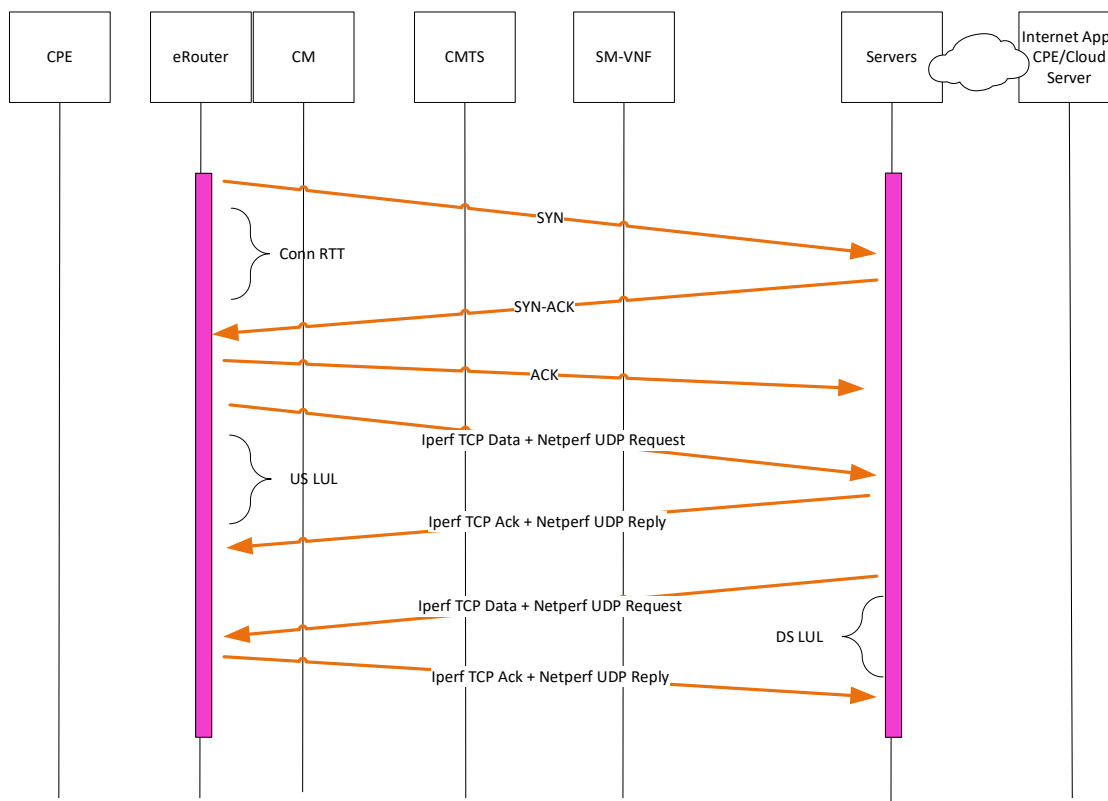


Google Docs, etc.) can open between 15-20 TCP connections. Delays in opening those TCP connections impact overall application performance, and thus impact the customer's overall experience.

While LM-1 helps the cable operators to assess the impact of typical home and serving group utilization levels on the network latency and jitter, LM-2 helps to assess the bufferbloating issues by measuring the latency and jitter at the highest home utilization levels.

The LM-2 method can be used within the MSO's network without requiring any HW test client in the subscriber's home as shown in Figure 3. In this case the test client is implemented within the eRouter firmware (e.g. RDK-B). The test can be initiated only when the subscriber's gateway (i.e. erouter+CM) has low utilization level not to degrade the subscriber's network performance. Hence, the connection RTT portion will be mostly ~10-20ms for current DOCSIS networks. This RTT includes media access delay in the US but queuing delay may not be high since the test is initiated when the home traffic volume is low.

Queuing delays in the CM and CMTS are the dominant factors in the US and DS LUL measurements respectively. UDP pings are impacted by queue building iperf TCP traffic since they are transmitted within the same service flow. LUL provides information on the latency variation subscribers may perceive when they use low latency applications. It reflects the impact of instantaneous bursts of queue building traffic (e.g. file download/upload, streaming) on the performance of the LL applications (e.g. gaming, videoconferencing).



**Figure 3 – Latency Under Load Measurement**

### 3. Latency Performance in Current and Emerging Network Architectures

#### DOCSIS NETWORKS LATENCY

As summarized in Table 2 – DOCSIS Networks Latency: RTT Between CM and CMTS Table 2 [3], the latency performance of DOCSIS networks has been improving with each version. A major delay source is queuing [3], that has been managed by the buffer control (BC) and active queue management (AQM) introduced in the D3.0 and D3.1 specifications. However, these improvements are not adequate for low latency service requirements and for providing MSOs a competitive edge in 10G era. Note that although CMs have mandatory AQM specifications (i.e. the PIE algorithm), queue management at the CMTS varies based on the vendor and its firmware releases. Table 2 levels correspond to latency under US load. As a result, latency under DS load and bi-directional load may be higher than shown in Table 2 for existing implementations. For example, we can confirm US LUL results with ~10-20ms for ~90<sup>th</sup> percentile, but when we run a bi-directional load, and depending on the CMTS implementation and configuration, the US LUL results can be in the order of a few hundred of milliseconds. CMTSs may have large physical queues. If cable operators do not optimize the configurations for the BC and AQM features of older CMTS deployments, they may create high queueing delays and inefficient network utilization. In Figure 4, below, we show some examples with test cases and configurations to depict the performance changes.

**Table 2 – DOCSIS Networks Latency: RTT Between CM and CMTS**

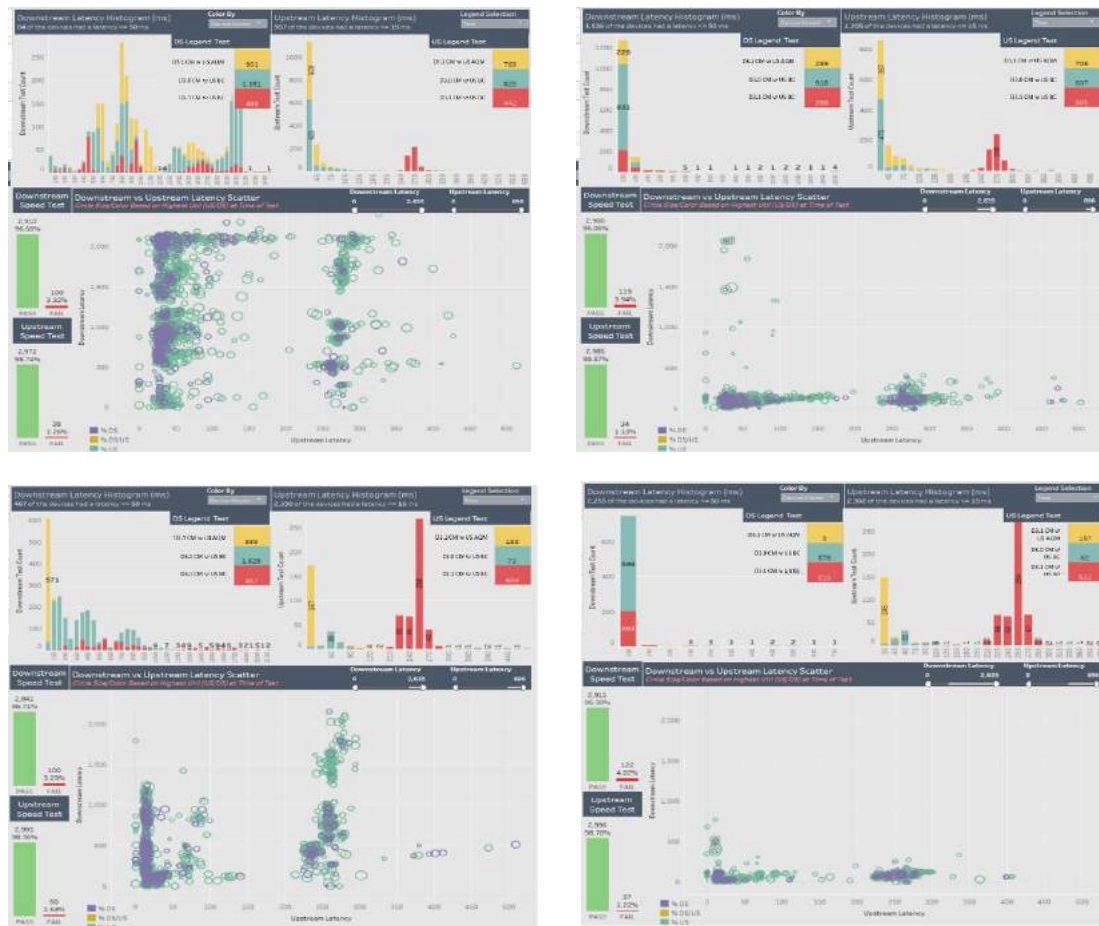
|                                       |            | When Idle | Under Load | 95-99 <sup>th</sup> Percentile |
|---------------------------------------|------------|-----------|------------|--------------------------------|
| DOCSIS 3.0 Early Equipment            |            | ~10ms     | ~1000ms    | ~1000ms                        |
| DOCSIS 3.0 w/ Buffer Control          |            | ~10ms     | ~100ms     | ~100ms                         |
| DOCSIS 3.1 w/ Active Queue Management |            | ~10ms     | ~10ms*     | ~100ms                         |
| Low Latency DOCSIS 3.1                | Dual Queue | <10ms     | <10ms      | <10ms                          |
|                                       | PGS        | ~1ms      | ~1ms       | ~1ms                           |

- Latency under DS load and bi-directional load may be ~50-100ms with current CMTS AQM implementations.

In the examples shown in Figure 4, the pre- and post-optimization settings for the DS AQM are shown. As seen in Figure 4, ~1-2 seconds of DS LUL is observed for pre-optimization, while post-optimization reduces this latency to the order of ~100-200ms. The suboptimal settings cause higher queueing delay caused by bursty queue-building traffic, as CMTSs may have large physical queues. Latency is reduced with the optimized settings and speed tests results show that there is no degradation in throughput. Note that outliers, e.g. some unexpected high latency results, due to test failures are not removed in these examples. When a D3.1 CM's US AQM is disabled, the default BC is 250ms as defined in the earlier D3.1 specifications, with US LUL around 250-300ms, which can be seen in the graphs (red CMs). This is also displayed in Figure 5 that shows the test US LUL results per speed tier rates and CM models, which are configured for different BC and AQM settings. The results show that when appropriate features are not enabled or configured with optimized settings, high US LUL may be observed. For example, a D3.1 CM with US AQM disabled and default BC settings (i.e. 250ms) has higher US LUL compared to D3.1CMs with US AQM enabled and D3.0 CMs with lower buffer sizes.

Figure 6 displays dslreports.com test results for the same CMTS and CM with and without optimized AQM settings. In this case, the test is done from a laptop connected via wired Ethernet to the CM. Optimized AQM settings improve LUL results (grade for bufferbloat) while still meeting throughput requirements.

These examples show that cable operators should measure LUL and audit CM AQM and BC configurations to make sure optimized settings are deployed.



**Figure 4 – Top Left: Max DS and US LUL with suboptimal DS AQM settings; Bottom Left: Mean DS and US LUL with suboptimal DS AQM settings; Top Right: Max DS and US LUL with optimized DS AQM settings; Bottom Right: Mean DS and US LUL with optimized DS AQM settings**


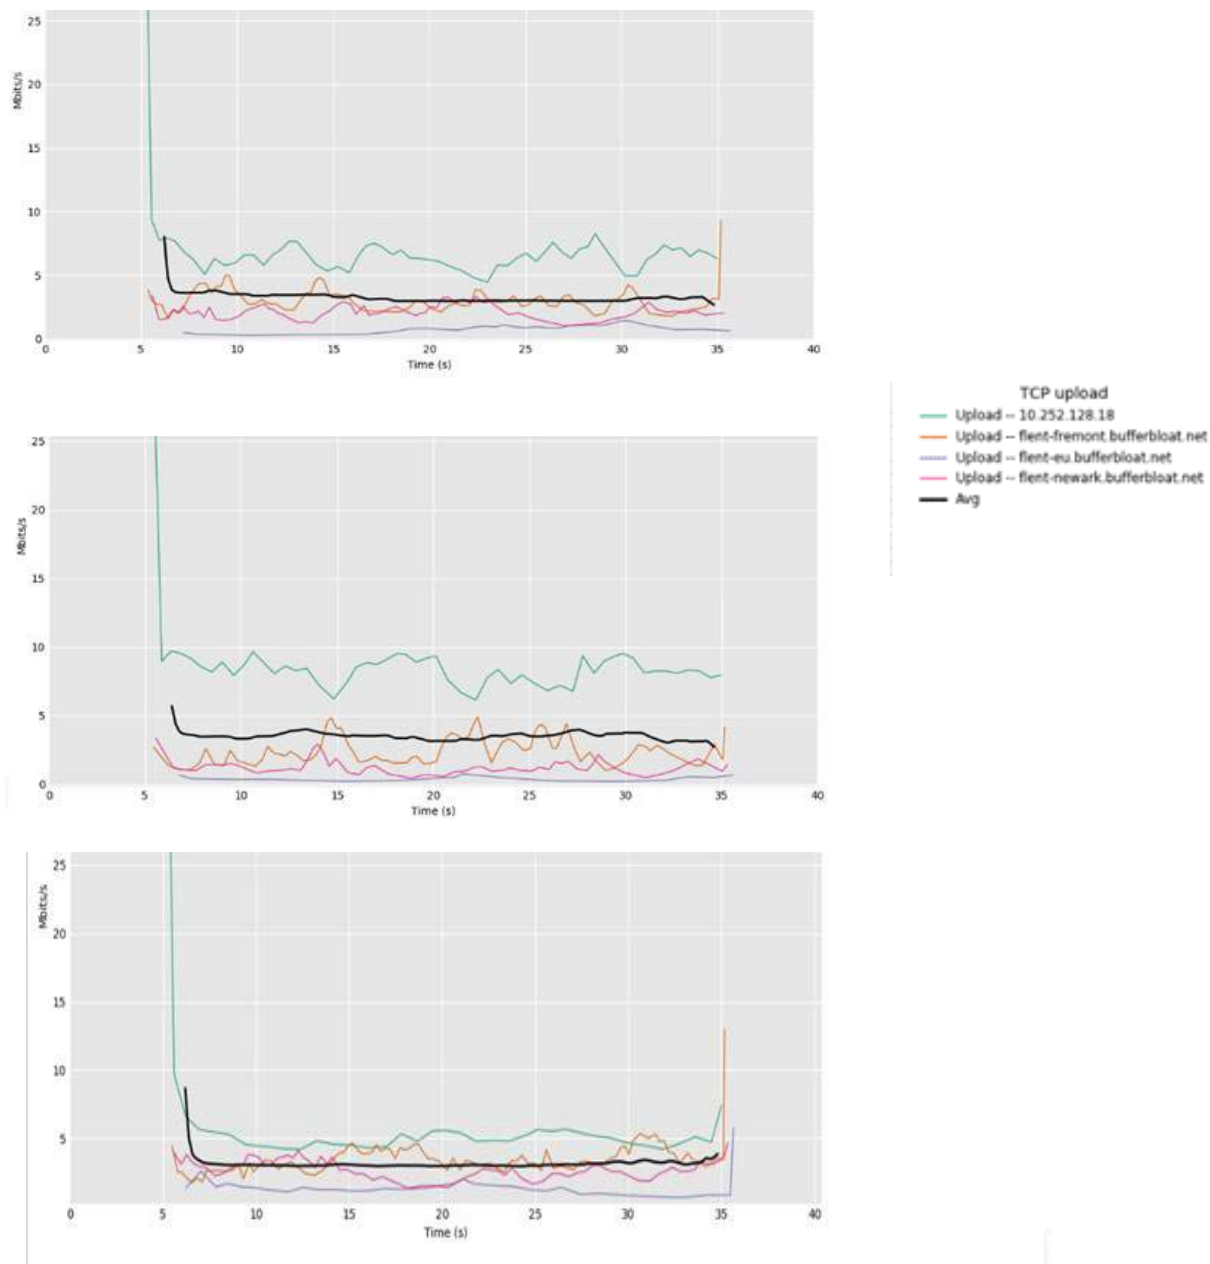


Figure 5 – Top: Max US LU; Bottom: Mean US LUL for different CM models and speed tiers



**Figure 6 – Left: dsreports.com results with suboptimal DS AQM settings and D3.1 CM with BC with 250ms default target latency; Right: Optimized DS AQM settings and D3.1 CM with US AQM with 10ms default target latency**

Depending on the buffer control and AQM, other factors, such as DOCSIS configurations, speed tier rates, additional RTT over the data path (e.g. additional delay at the home or northbound of the MSO's network), transport and congestion control protocols, duration of flows, rate adaptation schemes and the mix of traffic and utilization levels, all can affect the system performance. For example, Figure 7 shows that buffer sizes in a D3.0 CM affect how much throughput each concurrent flow within the subscriber's HSD SF gets, with different end-to-end RTTs. The graphs show US throughput for each flow within an HSD service of 10Mbps US speed tier rate. Although it may be desirable for a flow with a smaller RTT (e.g. an edge computing service) to have better performance, starving flows with long end-to-end RTTs should be avoided. In this example, 10, 30 and 50 ms target buffer sizes are set for the BC. 4 flows are destined to different servers, green flow's server is the closest to the subscriber's home while the purple flow's server is the farthest away. MSOs may improve their network with new features while optimizing configurations for the earlier versions of DOCSIS components.



**Figure 7 –Throughput (Mbps) values over time for each flow per CM HSD buffer size. Each flow has different e2e RTT. Green flow’s server is the closest to the subscriber’s home while purple flow’s server is the farthest away. Top: 10ms target buffer size; Middle: 30ms target buffer size; Bottom: 50ms target buffer size**

As described in Section 3, new D3.1 Low Latency DOCSIS (LLD) features with dual queue AQM and reduced bandwidth allocation map (MAP) interval time will enable non-queue-building low latency traffic to have an RTT of <10ms for the 95-99<sup>th</sup> percentile. This will improve gaming experiences, because gaming control traffic can be classified as NQB LL traffic and can be transmitted with much better jitter characteristics. Work is ongoing as it relates to scalable congestion control algorithms, so that streaming services can be also transmitted as NQB LL traffic. This area requires more research and development to analyze traffic characteristics and their impact on other service flows.

The LLD architecture is premised on the concept of separating Queue Building (QB) traffic from Non-Queue Building (NQB) traffic. An example of QB traffic is the typical file transfer. Online gaming control traffic is an example of NQB traffic. Current network deployments carry both QB & NQB traffic, transiting in the same service flow. When the service flow is unsaturated, this is not a problem. But when the service flow is at capacity, it is quite possible that NQB traffic gets stuck behind QB traffic. When this occurs, the NQB traffic incurs latency delays while the queue drains out the QB traffic and the NQB traffic awaits its transmit opportunity.

The LLD architecture attempts to alleviate this issue by separating QB & NQB traffic into separate sub-service flows, each with their own queue, as part of an Aggregate Service Flow (ASF). The available bandwidth of the overall ASF is still the same as what the customer has purchased, but the NQB & QB queues are drained in a manner that allows NQB traffic to be rapidly dispatch from the queue, while also ensuring that the QB queue receives adequate transmit opportunities.

Proactive grant scheduling is another D3.1 LLD feature, targeting a ~1ms RTT for the 99<sup>th</sup> percentile. Network efficiency for PGS must be analyzed for possible use cases.

## **Wi-Fi NETWORKS LATENCY**

Although avoiding Wi-Fi can result in more deterministic latency, jitter and packet loss, subscribers run low latency services such as gaming and videoconferencing over Wi-Fi all the time, because it is simply more convenient than wired Ethernet cables. Features such as WMM and AQM [4],[5] aim to improve latency and jitter for low latency services in Wi-Fi networks. Similar to DOCSIS networks, the tradeoffs between latency/jitter and throughput and fairness must be balanced (e.g. the tradeoff between efficiency with frame aggregation, vs faster channel access).

Although significant improvements may be observed with correct queue management and channel access control, factors such as outside interference and high concurrent utilization have been limiting the Wi-Fi performance for low latency services. The need for a more deterministic quality of service has been supported by Wi-Fi 6 (802.11ax) with MU-MIMO and OFDMA, enabling both DS and US increased simultaneous communications [7]. Removing contention along with deterministic QoS features (e.g. multi-user EDCA) that can provide low latency, jitter and overhead will be key to supporting low latency services.

In addition, Wi-Fi 6E, with its additional 1200 MHz in the 6 GHz spectral range, would be suitable especially for MDU type environments, but can be also used with multiple routers in large houses in the



future. Although 6 GHz is used by other networks, and more testing is required for range-speed-latency characterization, all of these improvements point to an optimistic future for low latency services. Most residential applications will not need Gbps symmetrical throughput anytime soon, but having this kind of speed both on the Wi-Fi and DOCSIS networks enables new services, including commercial and industrial services. Enterprise trials [6] have proven 2 Gbps speeds with consistent connections and latency around 2ms.

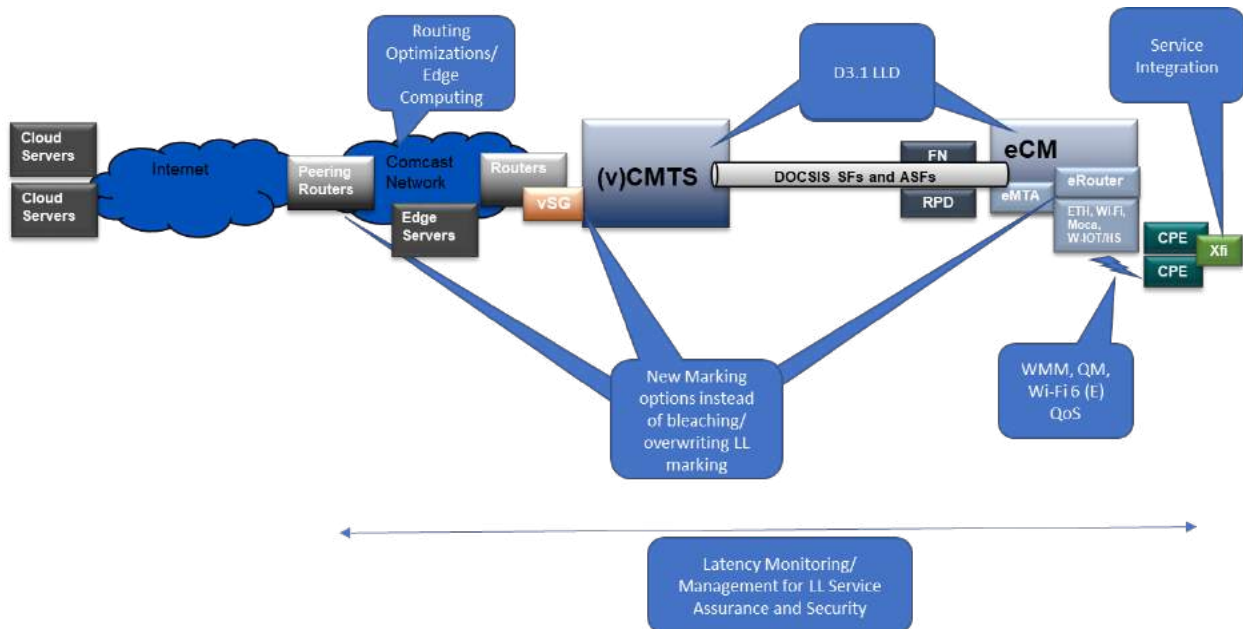
## CORE NETWORKS LATENCY

Game developers and providers, gaming router developers and network optimization companies have been developing products to optimized routing/tunneling to best gaming servers by measuring RTTs and/or establishing private networks. SDN-WAN has been applied for low latency services to intelligently shift traffic and dynamically adjust to network and traffic changes. Delivery at the edge by linking up data centers with ISPs' last mile networks has been proven to reduce latency and jitter significantly and improve fairness among online gamers.

MSOs have a unique position to provide the best end-to-end performance for gamers by applying low latency features at the access and home networks as well as at their edge core and peer routing platforms.

## 4. End-to-end Support For Low Latency Services

In addition to support LL service performance requirements, MSOs need to provide architecture changes for LL service classification, marking, service integration and assurance. Table 3 summarizes the main features for those architecture changes, and Figure 8 illustrates them for the corresponding components.



**Figure 8 – End-to-end LL services support**

**Table 3 – Architecture changes to support LL services**

| Current Architecture                                                                                                                                  | Target Architecture                                                                                                                     |
|-------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>• No classification and marking in untrusted network segments</li> </ul>                                       | <ul style="list-style-type: none"> <li>• End-to-end classification with new security measures</li> </ul>                                |
| <ul style="list-style-type: none"> <li>• Single queue for all HSD applications and scheduler (in DOCSIS and most Wi-Fi networks)</li> </ul>           | <ul style="list-style-type: none"> <li>• Dual/multiple queue and weighted scheduler for NQB-LL and other applications</li> </ul>        |
| <ul style="list-style-type: none"> <li>• Lower US speeds; Wi-Fi contention and unreliable QoS</li> </ul>                                              | <ul style="list-style-type: none"> <li>• &gt; Gbps DOCSIS and Wi-Fi with deterministic latency/jitter for LL services</li> </ul>        |
| <ul style="list-style-type: none"> <li>• Single performance measurement for all services within HSD SF</li> </ul>                                     | <ul style="list-style-type: none"> <li>• Differentiated performance measurement for NQB-LL and other applications</li> </ul>            |
| <ul style="list-style-type: none"> <li>• Single network optimization technique for all services within HSD SF in DOCSIS and Wi-Fi networks</li> </ul> | <ul style="list-style-type: none"> <li>• Network optimization techniques per traffic requirements within HSD service flows</li> </ul>   |
| <ul style="list-style-type: none"> <li>• Loose coupling in design, development, testing, and operations among network segments</li> </ul>             | <ul style="list-style-type: none"> <li>• Orchestrated end-to-end design, development, testing and operations for LL services</li> </ul> |
| <ul style="list-style-type: none"> <li>• Limited business models for HSD apps</li> </ul>                                                              | <ul style="list-style-type: none"> <li>• Flexible business models for LL services</li> </ul>                                            |

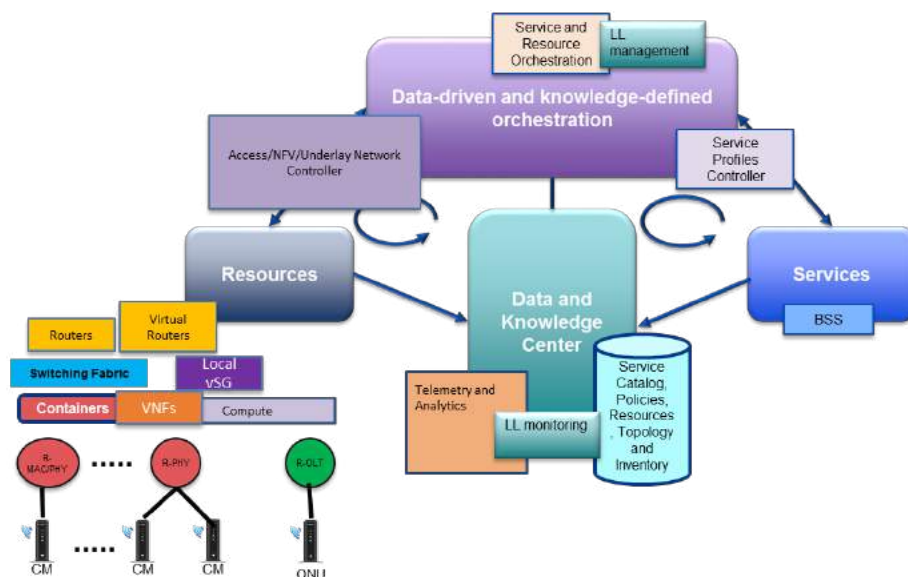
Marking options are being discussed in IETF WGs such as Low Latency, Low Loss, the Scalable (L4S) proposal and the Differentiated Services Code Point (DSCP) marking proposal[8][9]. Until these approaches are adopted widely, MSOs may support LL NQB traffic classification by changing their current marking options and taking new security measures. In this sense, another VNF envisioned for the VSG platform is in support of the Low Latency DOCSIS architecture.

This architecture can improve the customer experience by having QB & NQB traffic serviced, separately as addressed in Section 3. The challenge for the network operator then becomes how to separately identify NQB from QB traffic as it traverses the network. Assuming that business rules provide guidance on what qualifies as QB & NQB traffic, the data packets from each type of traffic must have some unique identifier, so that the CMTS can direct each type of traffic into the correct sub-service flow within the ASF. Here, a VNF can be deployed onto the VSG which includes rule mapping, to re-mark the data packets of the two types of traffic. This could be through changing the packet's DSCP value, based on other packet values, such as source or destination IP address, port number, protocol number, etc. A VNF supporting this re-marking can be dynamically updated to support changing business rules on the classification of QB vs NQB traffic, and re-mark (or not) accordingly. Thus, by providing a platform to instantiate this VNF, the VSG platform can play a part in implementing the overall LLD architecture live on the network.

VNFs such as a VSG platform are also crucial for end-to-end latency management systems. Figure 9 illustrates an example monitoring tool that collects hop-by-hop information such as network (including AQM, BC, and QoS MIBs), utilization and latency test results. A VSG or similar VNF can detect anomalies or changes that require correction actions. This information can also be used for performance



Figure 11 displays an architecture where resources may be managed and configured based on monitoring and prediction systems control [2]. LL monitoring and management can then be part of the data center as an orchestration function. Today, even some simple configurations may not be available, because of a lack of provisioning flexibility. For example, buffer sizes in older modems may be expressed in bytes, and setting them based on speed tiers may require the coordination of service class names and boot file settings. There may be other configuration parameters that need to be coordinated as well. Flexible configurations and operations may seem like features for the farther future, but today MSOs already deploy SDN/NFV enabled distributed systems and new telemetry platforms. Configuration flexibility is already an integral part of such systems. The next section discusses new architectures where low latency services may be integrated.



**Figure 11 – Low Latency Services Monitoring and Management Integrated within Data-driven and Knowledge-defined Architectures**

## 5. Conclusion: Final Thoughts on Latency Management

Although service assurance with performance management has been always the main driving force in designing access network architectures, a unified platform with an end-to-end orchestration approach has not been fully adopted, largely because of design limitations in many operators' networks. Recent changes in MSO network and service architectures provide the building blocks for such a unified platform, including:

- *Networking improvements:* New features in both wireline (e.g. DOCSIS/PON) and wireless (e.g. Wi-Fi) networking improve both the customer experience and network efficiency significantly. These features include frequency split and extended spectrum in DOCSIS, FDX, higher rate PON technologies and coherent optics, 802.11ax (Wi-Fi 6), low latency and distributed architectures for DOCSIS, Wi-Fi and mobile networks. Although some of these technologies aim at a specific hop or segment of the MSO network, initiatives like low latency networks target end-to-end improvements.
- *Data-driven networks and Monitoring:* Changes in telemetry, e.g. adapting push-based and cloud-hosted telemetry, enable data-driven network and service architectures. Both hop-by-hop, end-to-end latency and other performance metrics can be collected for overall latency management, troubleshooting, operations and planning purposes. Low latency services such as gaming will benefit from different latency measurement approaches, including concurrent and multi-hop measurements.
- *Software Defined Networking:* SDN enables centralized orchestration and coordination of the distributed controllers in different network and service segments. Dynamic and flexible configurations will help low latency services such as gaming, for example, by avoiding separate configurations of home and access network components. Some of the traditional latency measurement techniques assume that the control and data planes overlap, which wouldn't be the case for SDN networks. On the other hand, SDN enables an end-to-end data path view, with associated capabilities and monitoring that may be easily controlled for hop-by-hop and end-to-end measurements.
- *Network Function Virtualization:* MSOs have been introducing new VNFs over the control and data paths, with innovative functionalities in the areas of subscriber and service flow management that can help the differentiation of services per their traffic requirements. Virtualization in access networks may help to integrate new queueing and scheduling functionalities -- while special design requirements need to be considered for low latency services, as these designs may introduce additional latency not found in purpose-built, hardware-based architectures.
- *Knowledge-defined Networking:* Advances in the application of machine learning (ML) techniques to MSO networks and services open new doors for better performance prediction and management with self-optimizing capabilities. Recent advances in proactive network management (PNM) in DOCSIS and Wi-Fi networks can be extended for low latency services. The advantage of a knowledge-defined network is the ability to apply multi-hop PNM for end-to-end service assurance. In addition to smart networks and operations, MSOs have been using knowledge-defined systems for smart homes and customer interfacing platforms, which facilitate new low latency service offerings while assessing the customer experience for these services.
- *Cloud-based applications vs edge computing:* MSO integration of both cloud and edge computing based applications, depending on the service requirements, will enable low latency service providers to select the best architecture for the optimized customer experience. For example, cloud-based game providers can optimize network peering or consider edge computing based on performance, hardware and cost requirements.
- *Open source products and standards:* Many standardization efforts like the low-latency work within CableLabs and IETF enable support in a larger ecosystem. MSO use of open source products, and flexible integration of third party services in their platform, ensure the compliancy with regulations and policies as well.
- *Security:* Security, in terms of network and customer privacy protection, has become a vital item to be integrated into the design, instead of a later add-on feature. Low latency services require new classification and marking design, which requires new security elements. A proactive design approach is pivotal to a secure end-to-end solution, even while each network segment can have its own security feature (e.g. queue protection in D3.1 LLD specs).

- *Accessibility*: Similar to security in proactive inclusions, accessibility was a core part of the initial architecture design process for MSOs, instead of a later add-on feature. Low latency services may include strategies that offer control options to subscribers. MSOs that have an established framework to incorporate accessibility requirements early in the design can easily integrate low latency services by meeting every subscriber's needs.

## Abbreviations

|         |                                                 |
|---------|-------------------------------------------------|
| AR      | Augmented reality                               |
| ASF     | Aggregate service flow                          |
| AQM     | Active Queue Management                         |
| BC      | Buffer Control                                  |
| CM      | Cable modem                                     |
| CMTS    | Cable modem termination system                  |
| COTS    | Commercial off-the-shelf                        |
| DOCSIS  | Data over cable service interface specification |
| DS      | Downstream                                      |
| DSCP    | Differentiated Services Code Point              |
| IETF    | Internet Engineering Task Force                 |
| ISBE    | International Society of Broadband Experts      |
| LL      | Low Latency                                     |
| LLD     | Low Latency DOCSIS                              |
| LM      | Latency Measurement                             |
| LUL     | Latency Under Load                              |
| MAP     | Bandwidth Allocation MAP                        |
| ML      | Machine learning                                |
| MIB     | Management Information Base                     |
| MU-MIMO | Multi-User Multi-Input Multi-Output             |
| NFV     | Network Functions Virtualization                |
| NQB     | Non-queue-building                              |
| PNM     | Proactive Network Management                    |
| QoE     | Quality of experience                           |
| QoS     | Quality of service                              |
| QB      | Queue-building                                  |
| PIE     | Proportional Integral Enhanced                  |
| RTT     | Round-trip time                                 |
| SCTE    | Society of Cable Telecommunications Engineers   |
| TCP     | Transmission Control Protocol                   |
| SDN     | Software Defined Networking                     |
| US      | Upstream                                        |
| VNF     | Virtual Network Function                        |
| VR      | Virtual Reality                                 |
| VSG     | Virtual Subscriber Gateway                      |
| WG      | Working group                                   |
| WMM     | Wi-Fi MultiMedia                                |

## Bibliography & References

- [1] *Supporting The Changing Requirements For Online Gaming*, K. Scott Helms, SCTE-ISBE Workshop 2018
- [2] *The Future of Operations: Building a Data-Driven Strategy*, Sebnem Ozer, Sinan Onder, and Nagesh Nandiraju, “SCTE Journal on Network Operations, 2018
- [3] *Low Latency DOCSIS: Overview And Performance Characteristics*, Greg White, Karthik Sundaresan and Bob Briscoe, SCTE-ISBE Workshop 2019.
- [4] <https://www.cablelabs.com/10g/latency#:~:text=Our%20Low%20Latency%20DOCSIS%20technology,slowing%20down%20all%20other%20data.>
- [5] [https://www.bufferbloat.net/projects/bloat/wiki/What\\_can\\_I\\_do\\_about\\_Bufferbloat/](https://www.bufferbloat.net/projects/bloat/wiki/What_can_I_do_about_Bufferbloat/)
- [6] <https://wballiance.com/wbas-first-phase-of-wi-fi-6e-trials-shows-the-massive-potential-of-wi-fi-in-the-6ghz-band/>
- [7] *The Importance of Wi-Fi 6 Technology For Delivery of Gbps Internet Service*, David John Urban, SCTE-ISBE Workshop 2019.
- [8] <https://datatracker.ietf.org/doc/draft-ietf-tsvwg-l4s-arch/>
- [9] <https://datatracker.ietf.org/doc/draft-ietf-tsvwg-nqb/>

## Acknowledgments

The authors would like to thank and acknowledge all those who helped to make this paper possible. This paper includes several tests done and dashboards created by Aaron Tunstall, Sarulatha Subbaraj, Soomin Cho, Peifong Ren, Ray Hammer, Lei Zhou and Joe McHale.

# **Software Revolution Of Field Meters Using a Field-Capable Measurement Device**

A Technical Paper prepared for SCTE•ISBE by

**Anthony Curran**

Principal Eng II, Software Dev & Eng  
Comcast  
175 Washington St Norwell, MA 02061  
(781) 206-1389  
Anthony\_Curran@cable.comcast.com

**Andy Martushev**

Principal Eng II, Software Dev & Eng  
Comcast  
1401 Wynkoop St Suite 300 Denver, CO 80202  
(720) 512-3669  
Andy\_Martushev@cable.comcast.com



# Table of Contents

| Title                                         | Page Number |
|-----------------------------------------------|-------------|
| 1. Introduction .....                         | 3           |
| 2. Historic Context.....                      | 3           |
| 3. Problem / Opportunity Statement .....      | 4           |
| 4. Solution .....                             | 5           |
| 4.1. DOCSIS Based Cable Modem Gateway .....   | 5           |
| 4.2. Power Source.....                        | 5           |
| 4.3. Graphical User Interface .....           | 6           |
| 4.4. Performance Requirements .....           | 7           |
| 5. Operator Benefits.....                     | 8           |
| 5.1. Adoption Rate.....                       | 8           |
| 5.2. Employee Satisfaction .....              | 9           |
| 5.3. Ecosystem .....                          | 9           |
| 5.4. Cloud-Capable .....                      | 10          |
| 5.5. Immediate Historic Data Collection ..... | 10          |
| 5.6. Peer Assistance .....                    | 11          |
| 5.7. Training .....                           | 11          |
| 5.8. Real time analysis and feedback.....     | 11          |
| 5.9. Metrics .....                            | 12          |
| 5.10. Remote Diagnostics .....                | 12          |
| 5.11. Software Architecture .....             | 12          |
| 5.12. Reduce Development Costs .....          | 13          |
| 5.13. COVID-19 Opportunities.....             | 13          |
| 6. Field Feedback .....                       | 13          |
| 7. The Future .....                           | 14          |
| 8. Conclusion .....                           | 15          |
| 9. Acknowledgements .....                     | 15          |
| Abbreviations.....                            | 15          |
| Bibliography & References .....               | 16          |

## List of Figures

| Title                                                     | Page Number |
|-----------------------------------------------------------|-------------|
| Figure 1 – Jerrold 727 Field Strength Meter .....         | 4           |
| Figure 2 – Ingress Widget on phone .....                  | 6           |
| Figure 3 – Ingress Widget on Small Tablet.....            | 6           |
| Figure 4 – Ingress Widget on Large Tablet.....            | 7           |
| Figure 5 – Spectrum Widget on Desktop Computer.....       | 7           |
| Figure 6 – 1 <sup>st</sup> Generation Adoption Rate ..... | 9           |
| Figure 7 – 2 <sup>nd</sup> Generation Adoption Rate ..... | 9           |
| Figure 8 – Pre-Equalization Data Collection .....         | 10          |
| Figure 9 – OFDM Data Collection .....                     | 10          |
| Figure 10 – QAM Data Collection.....                      | 11          |
| Figure 11 – Peer Assistance Example .....                 | 11          |

# 1. Introduction

Field meters have been on the first line of defense when determining if a cable plant is sufficient for installs, in need of repair, and for determining the quality of service. The original meters were expensive, heavy, and required a large amount of training. Contemporary meters are still relatively high-cost and in some cases require even more training to use than their predecessors. In order to get a tool in every technician's hands, the device needs to be light, lower cost, and use contemporary software to display and interpret the relevant data to help reduce the amount of training required.

In the past 20 years, improvements in cable technology have occurred often and at a larger scale. From the late 1940s through the mid-1990s, most cable plant was used to retransmit analog TV signals. This represents 45+ years of similar technology, with most advancements involving increasing the upper spectral boundary of the infrastructure to support a greater channel capacity. By contrast, and with the adoption of DOCSIS 1.0 (1997), 2.0 (2002), 3.0 (2006), and 3.1 (2013), the cable industry has experienced a major technology upgrade approximately every five years. With each upgrade in cable technology, testing devices needed to evolve in parallel, to keep technicians current and to provide more relevant information to perform their duties.

With the onset of COVID-19, which impacted both field upgrades and social guidelines, tools needed to evolve even more rapidly. New safety protocols for cable install and repair technicians resulted in new signal verification and testing processes. Recent increases in service bandwidth needs are forcing quicker deployments of multiple orthogonal frequency division multiplexing (OFDM) and orthogonal frequency division multiple access (OFDMA) channels. By developing a software-first, cloud-connected device, the operator can adapt to the current situation, access more technology and evolve it more quickly.

This paper discusses how cable modem technology was used to fulfill this role, evolving to be part of the operator's ecosystem and bringing tremendous value. The device we developed and characterize in this paper is referred to as a "field-capable measurement device," or "FCMD."

## 2. Historic Context

Original cable signal meters were very large and cumbersome. Ron Hranac, Technical Marketing Engineer in Cisco's Cable Access Business Unit, would know. He started as a technician in Idaho in the 1970s, working with up to seven other technicians, all of whom shared Jerrold Model 727 field strength meters (they also used smaller installation-specific meters). This particular meter was 13" x 8.5" x 8.5" and weighed 15 lbs. It was also very expensive. So expensive that after one of them was dropped from a telephone pole, instead of buying a new one, Hranac once spent 40 hours repairing it back into usability.



**Figure 1 – Jerrold 727 Field Strength Meter**

Newer field meters are much smaller, measuring around 6" x 10" x 3" and weighing about four pounds. This is an improvement of approximately a factor of three in both size and weight, compared to those very early meters available to cable industry technicians. One added benefit of the current meters is that they are battery powered and usually have a runtime of a couple of days, depending on usage patterns.

Until the advent of DOCSIS, the method to measure cable signals primarily involved measuring analog radio frequency (RF) signals. The process required the technician to properly adjust gain factors for a specific channel, measure, then repeat this process for the next channel frequency. It was a very time consuming and labor intensive process.

The arrival of the cable digital age dramatically changed how cable modems were used. Steve Zanetich, Executive Director of Technical Operations at Comcast, has seen more than a few signal level meters in his career, dating back to the times of manually tuning each frequency, through adjusting a series of pots and knobs. In the digital age, those same meters jumped light years ahead of their predecessors in terms of features, processing, and measurement accuracy. "In the past, we needed to change hardware and software, recalibrate regularly, and [perform] other costly routine maintenance in an effort to keep measurements accurate," Zanetich noted. With today's technology, by contrast, considerably more can be accomplished, at a fraction of the cost. "The days of manually writing down your levels at each job are gone," he said. They've been replaced with real-time remote management, integrated business and performance tools, and collaborative software development, "all inside a feature-rich device that is capable of measuring so much more than just the amplitude of a single carrier."

### **3. Problem / Opportunity Statement**

Over the years, RF signal level meters (SLMs) have been made nearly indestructible, while being given the ability to make very accurate measurements. These requirements resulted in a meter with a significantly higher price point. The high price point, coupled with limited operator capital budgets, typically meant that most operators expected the meters to last between five and 10 years. With the rapid DOCSIS revisions, these meters can be obsolete before the cable operator can get a good return on investment (ROI).

The opportunity to use cable modems, which may not be as durable, but nonetheless make reliable and accurate measurements, was identified as a potential means to yield a better ROI.

After all, one element that is the same, between expensive spectrum analyzers and full-band-capable cable modems, is that they both thrive on the task that is demodulating signals. Developing a FCMD based upon a generic cable modem can help also achieve a good ROI.

## **4. Solution**

The solution was to design a FCMD tool based on of a portable cable modem, using known hardware commodities to help reduce overall cost. Similarly, the software architecture was selected from well-known open-source commodities, to again save costs in the short and long term. Environmental and technological performance requirements were adjusted to achieve a desired price point.

In the process of achieving these goals, numerous additional benefits were identified that had not been originally anticipated. These include periodic software improvements to keep the devices reasonably current; cloud connectivity to “share the load” of processing, and to reference historical or informational data too bulky to be stored locally; and the development of “role-specific” user interfaces (UIs) such as “line technicians” and “install and repair technicians”.

### **4.1. DOCSIS Based Cable Modem Gateway**

The first SLMs could be considered high-precision near-laboratory-grade instruments that were ruggedized to handle the harsh outdoor field environments we frequently encounter. With the advent of DOCSIS and the use of microprocessors, SLMs transitioned to be high-end pieces of field test equipment. All DOCSIS cable modems (CMs) are required to perform measurements/tasks similar to current field SLM equipment, just to function. Example of these measurements include:

- Full bandwidth spectrum capture
- Downstream signal level measurements for automatic gain control (AGC)
- Quadrature amplitude modulation (QAM) measurement error statistics
- Upstream equalization coefficients
- Upstream transmit power
- Data throughput (speed test)

Since cable modems perform tasks equivalent to meters, and a goal was/is to reduce the overall cost per device, the decision was made to develop the field-capable measurement device using a known commodity. As a direct result, a DOCSIS 3.0-capable (to be followed by DOCSIS 3.1-capable) cable modem was selected as the hardware platform.

The cable modem housing was designed consistent with its consumer premise equipment (CPE) version, to minimize ruggedization costs.

### **4.2. Power Source**

The development of a power source for the FCMD followed a similar methodology. Batteries have been used in our cable modem gateways for a long time, so the circuitry was extant. High capacity batteries have been around for a significant timeframe, and their technology is also proven. By selecting a generic form factor, we were able to decrease overall costs for the original product and any future battery replacements.

### 4.3. Graphical User Interface

Traditional SLMs are designed with a dedicated user interface as an integral component. This type of user interface usually touts benefits like being tightly coupled with the hardware, and ruggedized for daily field use. Those benefits can also represent drawbacks, however. Specifically, the drawbacks of an embedded graphical user interface (GUI) include:

- GUI is a possible point of failure requiring repair / replacement
- GUI software / GUI hardware is custom-designed on a per-product basis
- GUI hardware needs to be considered during power budget analysis
- GUI operating system includes licensing costs (sometimes but not always applicable)


At Comcast, technicians have been using existing network tools to interact with our cable modems for well over 10 years. Allowing them to interact with a portable cable modem, in the form of an FCMD, was an easy extension to those long-standing processes. The use of these tools brings additional benefits to the testing in the following ways:

- Cloud connectivity of the instrument (with consequent data ecosystem access)
- Common software development tools for application development
- Common operating system when mobile platforms are upgraded
- Global Positioning System (GPS) assistance with respect to application use
- Built-in, industry-standard accessibility features such as high contrast, large font sizes

Our technicians already use tools with adaptive and cross-platform form factors, suitable to their tasks, as depicted in Figures 2, 3, 4 and 5, which show the same ingress detection widget on a smart phone, small form factor tablet, large form factor tablet, and desktop computer, respectively:




**Figure 2 – Ingress Widget on phone**



**Figure 3 – Ingress Widget on Small Tablet**



**Figure 4 – Ingress Widget on Large Tablet**



**Figure 5 – Spectrum Widget on Desktop Computer**

#### **4.4. Performance Requirements**

Environmental requirements for high-end signal level meters are developed with the overall price of the product in mind. This means being able to withstand a variety of weather, environmental and electromagnetic interference (EMI) conditions. In addition, the RF front end of these meters is designed such that they have an extremely low noise figure, especially at low frequencies.

Based on over two years of field data, and despite fall-related ruggedization being a highly marketed feature of traditional SLMs, the truth of the matter is that product failures due to a drop are fewer than 0.1%. This is a very low calculated failure rate and was used as a consideration in the overall design. As a direct result, the FCMD was designed with the protection of an accompanying bag, designed to withstand a drop from a typical telephone pole height.

Environmental conditions, including rain, snow, heat and humidity, exemplify how cost reduction compromises carry some risks. For instance, temperature requirements required the addition of a fan on the second generation FCMD. The location of the ventilation holes, as well as the bag design, limited this risk. Over two years of data, the failure of this product due to environmental conditions is 0.2%. Again, with some mitigation in the design, the cost offset outweighed the risk.

Electrostatic discharge (ESD) and RF interference (RFI) issues are problematic for any cable field device. During initial development of the FCMD, multiple changes were made to minimize both internal and external issues which would affect performance.

High performance SLMs have an extremely low noise floor and a high dynamic range with respect to their input signal. A compromise was made during the design to accept and minimize the noise that is inversely proportional to the frequency common at low frequencies, 1/F noise. For a majority of technician tasks, the need is to see ingress about -30 decibel millivolt (dBmV) at 5 megahertz (MHz) – versus the ability to see ingress below 50 dBmV at 5 MHz. Specialized software was used to compensate for the front end design to achieve a sufficient dynamic range for ingress measurements, for both residential and line technicians.

## 5. Operator Benefits


The first generation FCMD device was based on DOCSIS 3.0 technology and was deployed in the fall of 2016. The second generation version was based on DOCSIS 3.1 and was deployed in fall of 2018. Beyond the lower cost, the benefits of the design were quickly realized over the past four years. By combining a non-proprietary software architecture with cloud connectivity, a number of features turned into benefits that were realized from this measurement device, versus a typical SLM. These feature/benefit pairings are bulleted below:

- *Reduced cost.* At a lower per-unit cost, the FCMD was and can be made available to more technicians, which translates into a higher adoption rate. It continues to increase employee satisfaction, and can be integrated into a larger tools and business metrics ecosystem within the company.
- *Cloud capable.* By linking the FCMD to cloud-based databases and processing functions, technicians gain immediate access to historic diagnostic data about a home or business, as well as performance metrics, remote diagnostics, and training opportunities.
- *Non-proprietary software architecture.* By taking an open-source, non-proprietary route to development, we were able to move in lockstep with agile software methodologies, which also reduced costs and improved employee satisfaction.


### 5.1. Adoption Rate

Perhaps not surprisingly, the capital expense budgets for network service providers are tightly controlled. By introducing a lower cost measurement device, operators can increase adoption rates, because more devices can be purchased from the same fixed capital budget. The first generation FCMD was, as previously mentioned, introduced in 2016, and achieved a 98% adoption rate across the operator's install and repair technicians within 13 months. This initial device was DOCSIS 3.0-capable, and despite a high adoption rate, its functionality didn't meet the need for a DOCSIS 3.1-capable device, especially for maintenance technicians.

The second generation FCMD device (DOCSIS 3.1-capable) was deployed in January of 2019. It achieved a 95% adoption rate for maintenance technicians within 16 months. Graphs showing the uptake for the first and second generation FCMD are shown in Figures 6 and 7.



**Figure 6 – 1<sup>st</sup> Generation Adoption Rate**



**Figure 7 – 2<sup>nd</sup> Generation Adoption Rate**

## 5.2. Employee Satisfaction

As mentioned previously, traditional SLM costs and functional lifetimes typically resulted in employees using the same meter for a minimum of five years. This can impact employee satisfaction, noted Tom Bach, Senior VP of Engineering for Comcast, who explained that employee morale can be affected, “especially when less-tenured employees are using the latest 3.0 DOCSIS-capable equipment, while long-term employees are still on DOCSIS 2.0-capable equipment.” The ability to equip all technicians with the latest DOCSIS 3.1 technology quickly helped to improve overall employee satisfaction, he added.

## 5.3. Ecosystem

With the low cost and consequent high adoption rate across the company, the FCMD product could be designed to participate in the company’s extensive and growing software ecosystem. All techs could be assigned identical devices, whether their primary roles involve installation and repair, maintenance, or headend tasks. This also allows data to be shared across functional groups. In turn, and over time, this improves the accountability and access to data across different populations of technicians.



The high adoption rate, coupled with being part of the company’s overall data-driven ecosystem, and the availability of the cloud, allowed the FCMD product to go beyond a field measurement device, to a device that supports training opportunities, metrics, and even remote diagnostic capabilities.

## 5.4. Cloud-Capable

Existing SLMs do, of course, have the ability to work with a cloud infrastructure. However, access to proprietary cloud data can be cumbersome and cost prohibitive. By contrast, the FCMD became part of the operator’s ecosystem, such that we were able to rapidly adapt its design to numerous cloud-capable applications.

Noted Larry Wolcott, Engineering Fellow at Comcast and a lead proponent of the industry’s proactive network maintenance (PNM) efforts: “Because of the open nature of the software APIs, we were able to quickly attach this tool into our software ecosystem. Other systems had clouds, but they were not adaptable into the rest of the ecosystem.”

## 5.5. Immediate Historic Data Collection

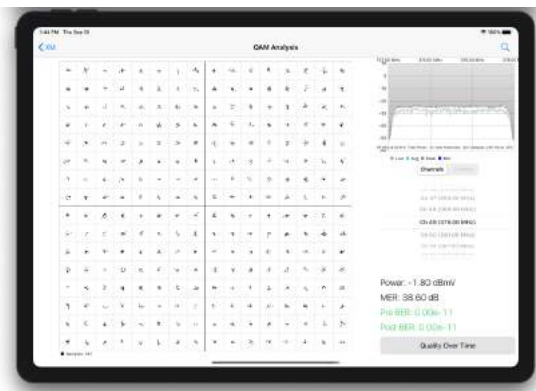
FCMDs have collected multiple terabytes of data over the past two years. The data collected include but are not limited to: ingress, full band spectrum capture, QAM channel signal analysis, OFDM signal analysis, and speed test performance. This information carries many long-term uses, including peer review, training, artificial intelligence (AI) and machine learning (ML) analysis, and system metrics. Figures 8, 9 and 10 depict different types of data collection using the FCMD:



Figure 8 – Pre-Equalization Data Collection



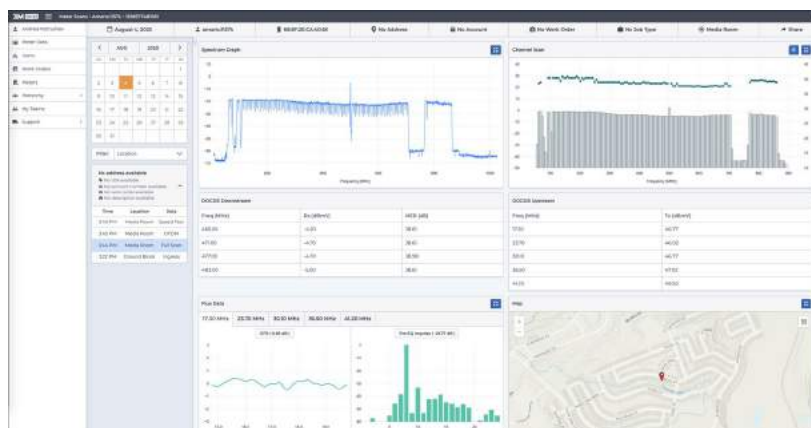
Figure 9 – OFDM Data Collection



**Figure 10 – QAM Data Collection**

## 5.6. Peer Assistance

Many field problems are difficult to diagnose even by an experienced technician. Using the technician's tools as a medium to share data immediately to the cloud allows other technicians and supervisors to assist technicians remotely. The need for multiple truck rolls to diagnose and repair a difficult problem can be and was decreased. Figure 11 shows an example of a peer-assisted diagnosis, where a remote technician can assist the technician that is onsite by reviewing signal spectrum, DOCSIS and video channel powers, and DOCSIS locked signals statistics.



**Figure 11 – Peer Assistance Example**

## 5.7. Training

Post-job analysis of historic data per technician provides an opportunity to design individualized training, per employee, to improve overall cable plant health in the short and long term. For example, supervisors routinely review employee ingress measurement metrics and work with each technician to identify what would likely cause the ingress seen on their jobs and possible solutions for future reference.

## 5.8. Real time analysis and feedback

Measurement data, at the time of collection from the FCMD, is flowed into data streams that serve multiple purposes: overall company consumption; real time analysis for network performance; and adherence to employee metrics. The data is also stored for post processing and long-term data mining.

## **5.9. Metrics**

For field technicians, metrics collections can simultaneously represent tremendous value and occasional contention – the latter because the addition of the cloud allows those metrics to be collected and validated on a per-job basis, which ultimately decreases the potential for obfuscation by the technician.

On the other hand, the ability to collect and analyze anonymized per-customer/per-incident performance data, over time, gives the operator the ability to document performance and determine possible customer improvements. These metrics help to improve the overall quality of the plant. In numerous cases, when we implemented a required set of metrics, a measurable decrease in repeat calls was seen.

Beyond per-technician or per-customer metrics used for diagnostics, the ability to collect meter application metrics has also been beneficial to identify the application usage statistics, which assist in application design. With the data available, the operator can understand which application widgets are being used and how much time it takes. Such metrics give opportunities to improve the widgets and employee training.

The collection of metrics into the ecosystem can also improve construction build-out processes. Upon completion of construction, the mobile measurement tool and ecosystem in the FCMD is able to document the RF performance of the new plant construction prior to operator acceptance – a capability deemed invaluable to building out plant infrastructure.

Metrics about various tasks can also be validated, by collecting demonstrated plant performance improvements that were observed on a job. One example might be when a speed test improves from 500 Mbps to 1 Gbps.

## **5.10. Remote Diagnostics**

The addition of the cloud and Wi-Fi connectivity also allows for remote control and remote diagnostics through the FCMD. A technician can now connect the device to the ground block, establish a DOCSIS connection, and, through a remote application, move around the premises, fixing issues and receiving immediate feedback from the device, still out by the ground block.

In addition, the cloud allows the FCMD to become a short-term piece of diagnostic equipment. Through the ecosystem, many of the technician's applications have been re-hosted in the cloud and can be controlled remotely.

## **5.11. Software Architecture**

A non-proprietary software architecture also enabled short-term product improvements that helped to improve overall acceptance and performance. Technicians have suggested numerous areas where the technology could be improved. Including technicians in the design feedback loop translates into quick product improvements. It also increases employee morale, by not only giving technicians a say in how they would like the gear to work, but also by rapidly implementing those suggestions.

“We get to make it do exactly what we want,” said Patrick Stephens, a Senior Business Technician at Comcast, who also noted that “it makes honest technicians.”

## 5.12. Reduce Development Costs

DOCSIS 3.0, DOCSIS 3.1, and high-split (5 MHz to 204 MHz upstream) cable modems have an underlying software architecture and operating system that is almost identical. Adding mid-split or high-split capability to an SLM is primarily a hardware (HW) upgrade and only requires minimal software upgrades. The jump to a DOCSIS 4.0-based cable modem will require additional software upgrades, but the backwards-compatible nature of DOCSIS indicates that a large amount of the existing development will be reusable.

Personal computer devices (desktops, laptops, tablets) continually evolve, yet the underlying operating system is usually 90% or more backwards-compatible. Similarly, a GUI that uses an operating system that has not significantly changed over the past five years can also dramatically decrease ongoing development costs. Additionally, with the mobile cable modem and control-plane in the FCMD being separate, either can be upgraded with minimal impact to the overall system.

## 5.13. COVID-19 Opportunities

Along the way, we've been able to identify and implement a quarterly improvement program to modify the FCMD's functionality, through software, to meet changing needs of the technicians.

One example of such dynamic adaptation, attributable to a non-proprietary software architecture and a cloud-based ecosystem, is the changes required to work in the physically distanced COVID-19 conditions that occurred in early 2020. Employee safety was crucial during this time, and required the development of new testing methodologies. Using existing and available field meters, processes would need to be identified and appropriate training would need to be developed and implemented. Additionally, there was no practical way to ensure that the process would be properly followed.

Our tools software architecture allowed us to build new signal processing algorithms to validate performance at taps, residential ground blocks and internal CPE -- without endangering employee safety by entering a premise containing potentially infected residents.

This is not the only use case where the software architecture has afforded us the ability to improve the product between hardware improvement cycles. We were able to rapidly build other tools, as well, to measure downstream alignment, improve speed test diagnostics, and log DOCSIS information.

## 6. Field Feedback

The integration of the SLM into our overall data ecosystem provided a number of time-saving and performance benefits. Todd Szuter, Director, Preventive Maintenance at Comcast, contributed this short-list of benefits:

- Old methods of collecting per-channel data on plotting paper, to see a standing wave, are gone.
- Instant cloud access dramatically decreased configuration headaches for residential technicians. No longer were techs using outdated channel maps, resulting in invalid refer to maintenance (RTM ) requests.
- Third party SLMs have cloud capabilities but are not integrated into our software ecosystem. The data may be recorded but it was rarely used.
- This cloud-enabled device keeps technicians honest by making data available. Those who want to succeed embrace the device.

- This capability has great value in validating new construction prior to marking the effort as complete. Performance metrics and pictures, including GPS data, can be recorded and validated prior to company acceptance
- Cloud metrics helped improve the plant by enabling all technicians to reduce issues at all jobs.

Szuter also identified some negatives and resistance with introducing new technology to the workforce. Many technicians, he noted, just don't want to switch, because embracing a new technology and/or paradigm can be difficult. Also, with the addition of cloud metrics, "even though we dramatically increased network health, reduced repeat rates and improved the overall customer experience, in some cases a job that used to take 10 minutes may take 50 minutes," he noted.

When asked about the challenges and rewards of pushing new technology so deeply into the field, Comcast's Zanetich, quoted earlier, reiterated that good design yields good usage. "I think with anything new, you will experience challenges, and the deployments of this platform were no different," he said. Specifically, encouraging technicians to embrace the idea of a portable FCMD, and trust its readings, takes time, but not that much time: "It didn't take long for our technicians to see the value of the powerful platform that we placed in their hands, which they themselves can iterate based on what they need," he continued. "It's amazing to see the changes in the field, that can now happen overnight with a software update – changes that used to require us to purchase a new device and spend months distributing."

Clearly, the adoption of new technology, especially test equipment, always comes with its challenges. With the introduction of the FCMD, technicians were able to find issues faster, improve rework metrics and optimize the customer experience. Gone are the days of hunting-and-pecking, from a troubleshooting perspective; as a result, technicians appreciate the platform and as a means to pinpoint what needs to be done to optimize product and services delivery. "As I talk to our technicians," Zanetich concluded, "they appreciate the investment the company is making to ensure they have the best equipment available."

## 7. The Future

When measurement tools are cloud-based, remote signal analysis can be done to reduce truck rolls – in some cases, truck rolls that weren't necessary in the first place. Eventually, the same software can be embedded in CPE like set-tops and gateways, because they all ultimately use the same DOCSIS underpinnings. The network becomes the test apparatus – and it runs constantly, finding problems, and self-healing when possible, via PNM, dynamic channel changes, and related adjustments: Informing customers so that they can help themselves, such as about loose connectors or too many splitters, or automatically creating outage information that can inform operations centers and network techs.

When asked for his opinion about adding more PNM features into CPE, Cisco's Hranac was enthusiastic. "I love it – the more information from premises equipment, the better, including what full band capture brought to cable modems, and PNM brought in DOCSIS 3.1-based cable modems."

Virtually all of the FCMD's functionality is based on software commonly found in all DOCSIS 3.1 compliant cable modems. The opportunity to add the underlying functionality to STBs and other CPE allows a large percentage of current truck roll analysis to be done via the cloud or autonomously. Currently, full band spectrum analysis is being performed on more than 85% of all deployed cable modems at Comcast. The ability to add constellation analysis on all channels beyond those representing the currently bonded channels, per modem, could yield additional insights about Moving Picture Experts Group (MPEG) video statistics, again to potentially prevent unnecessary truck rolls.

Meanwhile, DOCSIS 4.0 is coming. In order to be able to not limit the technology because of hardware limitations, the hardware should be overdesigned to allow options to be considered in software. As DOCSIS 4.0 evolves, there will likely be differences and improvements – and, as a result, any field tool should have the ability to accept over-the-air (OTA) updates to keep the devices in step with technology.

As a general rule, technician tools are continually getting smaller, faster, lighter and less expensive. Cable operators started with a field meter that was 25 pounds and required an external power source. Similar devices now are the size of a small loaf of bread. With the addition of software defined radios (SDRs) and system on a chip (SOC) designs, devices like cable modems and FCMDs at some point could be the size of a USB-connected Ethernet dongle.

The addition of AI/ML on all premise devices may improve troubleshooting. Cable modems and STBs can report network events at the premise level, between themselves. The aggregation of multiple events across devices sharing the same ground block can determine, for instance, if an issue is a single device issue, or one that affects the entire premise.

## 8. Conclusion

The software revolution of the FCMD has shown great benefits. The change to the ecosystem and returns on investment have so far outweighed the risks of this endeavor.

With the unanticipated adaptations to COVID-19 in the Spring of 2020, we demonstrated that the FCMD can quickly acclimatize to a new world, through software improvements and OTA updates, which brought tremendous value to the field.

This “cable modem as an analysis tool” paradigm, embodied in the FCMD, could apply to STBs and broadband gateways, as well. The ability to continuously innovate on this platform, depending on the needs of the operator, removes the hardware-based and proprietary constraints of the past. The software adaptability of the meter has accelerated innovation, improved employee and customer experiences, reduced cost and removed risk in our DOCSIS upgrades.

## 9. Acknowledgements

Much of the material within this document was sourced from interviews and written contributions of many Comcast leaders and industry experts. The authors would like to thank Ron Hranac from Cisco, and Tom Bach, Larry Wolcott, Stephen Zanetich, Todd Szuter, Rob Gonsalves, and Patrick Stephens from Comcast.

## Abbreviations

|        |                                                  |
|--------|--------------------------------------------------|
| AGC    | automatic gain control                           |
| AI     | artificial intelligence                          |
| CM     | cable modem                                      |
| CPE    | consumer premise equipment                       |
| dBmV   | decibel millivolt                                |
| DOCSIS | Data-Over-Cable Service Interface Specifications |
| EMI    | electromagnetic interference                     |

|       |                                               |
|-------|-----------------------------------------------|
| ESD   | electrostatic discharge                       |
| FCMD  | field-capable measurement device              |
| Gbps  | gigabits per second                           |
| GPS   | Global Positioning System                     |
| GUI   | graphical user interface                      |
| HW    | hardware                                      |
| Mbps  | megabits per second                           |
| MHz   | megahertz                                     |
| ML    | machine learning                              |
| OFDM  | orthogonal frequency division multiplexing    |
| OFDMA | orthogonal frequency division multiple access |
| OTA   | over-the-air                                  |
| PNM   | proactive network maintenance                 |
| QAM   | quadrature amplitude modulation               |
| RF    | radio frequency                               |
| RFI   | radio frequency interference                  |
| ROI   | return on investment                          |
| RTM   | refer to maintenance                          |
| SDR   | software defined radio                        |
| SLM   | signal level meter                            |
| SOC   | system on a chip                              |
| STB   | set-top box                                   |
| UI    | user interface                                |

## Bibliography & References

Ron Hranac, Technical Marketing Engineer in Cisco's Cable Access Business Unit

Larry Wolcott, Comcast Fellow

Tom Bach, Comcast SVP Engineering

Stephen Zanetich, Comcast Executive Director, Technical Operations

Todd Szuter, Comcast Director of Preventive Maintenance

Rob Gonsalves, Comcast Senior Director Product Development & Engineering

Patrick Stephens, Comcast Senior Business Technician

DOCSIS 3.0 Specification, Cablelabs.com

DOCSIS 3.1 Specification, Cablelabs.com

DOCSIS 4.0 Specification, Cablelabs.com

Radio Museum, <https://www.radiomuseum.org>

# Building the RPHY Upstream Scheduler with YANG

A Technical Paper prepared for SCTE•ISBE by

**Tong Liu, PhD**

Principal Engineer  
Cisco Systems Inc

300 Beaver Brook Road, BOXBOROUGH, MA 01719  
978-936-1217  
tonliu@cisco.com

**John T Chapman**

CTO Cable Access and Cisco Fellow  
Cisco Systems Inc

170 W Tasman Dr, San Jose, CA 92677  
408-526-7651  
jchapman@cisco.com



# Table of Contents

| Title                                                          | Page Number |
|----------------------------------------------------------------|-------------|
| 1. Introduction.....                                           | 4           |
| 2. Why Remote US Scheduler .....                               | 4           |
| 3. How to Spin off the Remote US Scheduler .....               | 5           |
| 4. Why YANG Model-Driven API for the Remote US Scheduler ..... | 7           |
| 4.1. A Formal Definition of Service Contract.....              | 7           |
| 4.2. Mature and Well Adopted.....                              | 8           |
| 4.3. Interoperation Made Easy .....                            | 8           |
| 4.4. Sufficent Performance at Scale .....                      | 8           |
| 4.5. Easy to Maintain.....                                     | 10          |
| 5. Remote Upstream Scheduler Behavior Model.....               | 10          |
| 6. YANG Model for Remote Upstream Scheduler .....              | 12          |
| 6.1. Base remote-us-scheduler YANG Module .....                | 13          |
| 6.2. us-scheduler-domain YANG Module .....                     | 13          |
| 6.3. us-qos-scheduler YANG Module.....                         | 14          |
| 6.1. Map-builder YANG Module .....                             | 15          |
| 7. Coexistence with Existing RPD Management Interface.....     | 17          |
| 8. Remote US Scheduler in DOCSIS YANG Echosystem.....          | 17          |
| 9. Conclusion.....                                             | 18          |
| Abbreviations .....                                            | 19          |
| Bibliography & References.....                                 | 19          |

## List of Figures

| Title                                                                           | Page Number |
|---------------------------------------------------------------------------------|-------------|
| Figure 1 – Centralized vs. Remote Upstream Scheduler Deployment Scenarios ..... | 5           |
| Figure 2 – Centralized Upstream Scheduler in R-PHY Today.....                   | 6           |
| Figure 3 – Remote Upstream Scheduler Architecture .....                         | 7           |
| Figure 4 – YANG Data Model-Driven Management Components .....                   | 8           |
| Figure 5 – Adding a UGS SF in a CM Initiated DSA Transaction.....               | 9           |
| Figure 6 – Adding a RNG-REQ Opportunity in MAP .....                            | 10          |
| Figure 7 – UGS Behavior Model .....                                             | 11          |
| Figure 8 – UGS Diagram and UML Model .....                                      | 12          |
| Figure 9 – remote-us-scheduler base module tree diagram and UML .....           | 13          |
| Figure 10 – us-scheduler-domain module top-three level tree diagram .....       | 13          |
| Figure 11 – us-qos-scheduler module top level tree diagram.....                 | 14          |
| Figure 12 – best-effort-scheduler tree diagram.....                             | 15          |
| Figure 13 – map-builder top level tree diagram .....                            | 16          |
| Figure 14 – Coexistence with Existing RPD Management Interface .....            | 17          |
| Figure 15 – CCAP Network Element Types and the Subclassing Hierarchy.....       | 18          |
| Figure 16 – Remote US Scheduler Module in DOCSIS YANG Echosystem .....          | 18          |

## List of Tables

| <b>Title</b>                                                     | <b>Page Number</b> |
|------------------------------------------------------------------|--------------------|
| Table 1 – US Scheduler Clients at the DOCSIS Control Plane ..... | 6                  |

# 1. Introduction

The remote upstream (US) scheduler is a low latency Remote PHY solution proposed in [1] to maintain low latency for DOCSIS regardless of the CIN distance. Moving the US scheduler from the CCAP core to the remote PHY device (RPD) creates a new management interface between the remote US scheduler at the RPD and the upper MAC layer clients at the CCAP core. This entails a robust and standard application programming interface (API) for easy configuration / management of multiple remote US schedulers at scale and enabling interoperability between different CCAP core and RPD vendors.

Towards this goal, we propose a YANG based API to manage the remote US scheduler. YANG is an API contract language widely used in the world of networking and is now being introduced into the world of DOCSIS. A specification written in YANG is referred as a “YANG module”, and a set of YANG modules are collectively called a “YANG model”. A YANG model characterizes the behavior of a network function with data hosted by the server that a client can manipulate and observe using standardized operations. Once the YANG model is published, both client and server can have faith that the other knows the syntax and semantics behind the modeled data.

In this paper, we describe how we used YANG to define the remote US scheduler service contract. Our YANG model has several modules. The first is the US scheduler itself that provides the scheduling services such as best effort, rtPS, nrtPS, UGS and PGS. The second component is the MAP Builder that allocates bandwidth across the US RF channels. The remote upstream YANG model is intended to be included in the CableLabs’ standard modules and become an integral part of the standard DOCSIS YANG ecosystem.

The rest of the paper is organized as follows. Section 2 explains why the remote US scheduler may help maintain low-latency for DOCSIS in R-PHY deployments. Section 3 shows how to separate the real-time scheduling functions from the CCAP core while keeping DOCSIS control plane intact. Section 4 examines the reasoning for using YANG data model-driven management for the remote US scheduler. Section 5 explains the fundamental modeling principles to establish the remote US scheduler YANG model. Section 6 presents the remote US scheduler YANG model structure and the key components. Section 7 discusses how to add the YANG based remote US scheduler to existing RPDs that are managed by the legacy Generic Control Protocol (GCP). Section 8 describes how to integrate the remote US scheduler into the upcoming DOCSIS YANG ecosystem. Finally, section 9 concludes the paper and highlights the takeaways.

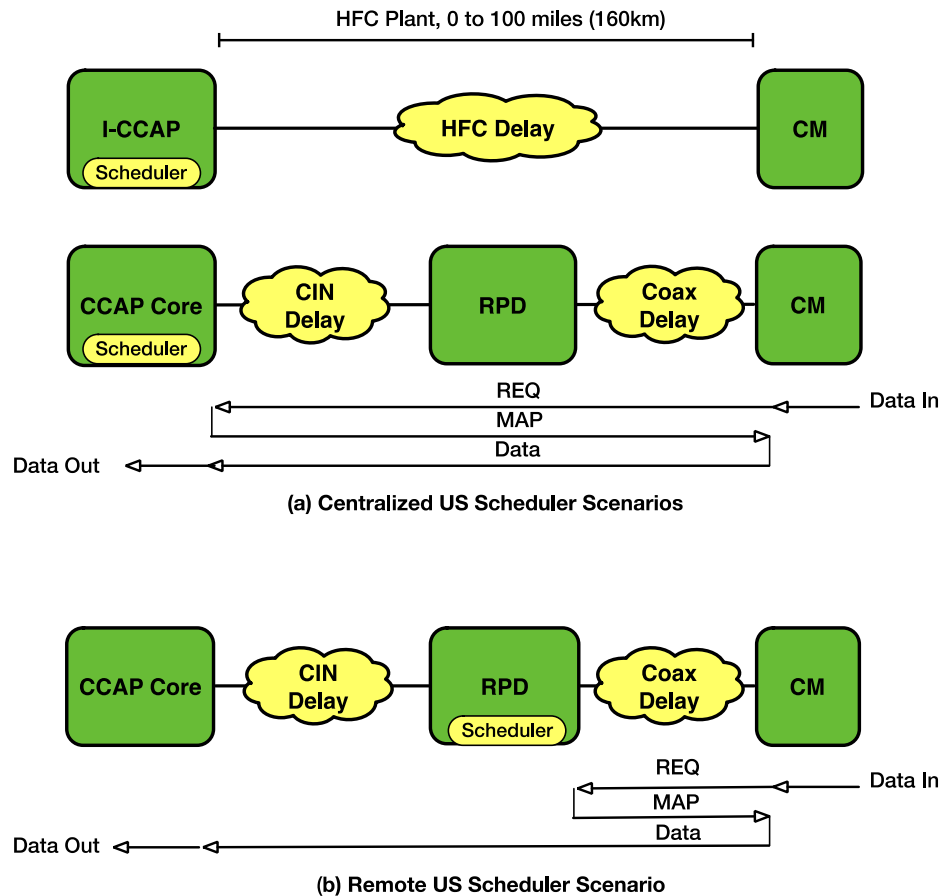
## 2. Why Remote US Scheduler

In the R-PHY architecture today, the US scheduler is centralized at the CCAP core, leaving only the PHY elements at the RPD. This design choice fits most of the R-PHY deployment cases today where the CIN distance is less than 100 miles and the MAPs run at 2-millisecond interval. Under these operation conditions, CIN delay has no impact on the upstream scheduling latency based on the analysis in [1], and R-PHY system is equivalent to the I-CCAP (Integrated Converged Cable Access Platform) in terms of the latency performance, but with much better RF capacity.

As the distributed access architecture (DAA) transformation deepens, there are scenarios where the CIN is stretched beyond the 100-mile mark, for reasons such as hub-side consolidation that relocates a CCAP core to the central headend or a regional data center. Meanwhile, driven by the new low latency applications, such as cloud gaming and mobile xHaul, the DOCSIS upstream request-grant (REQ-GNT) protocol is tightened to a shorter MAP interval at 1 millisecond [2]. Under these new operation conditions, the CIN becomes contributing factor to the upstream scheduling latency. By relocating the

upstream scheduler to the RPD, the CIN delay factor can be effectively removed from the REQ-GNT loop, therefore achieving latency improvement.

Figure 1 shows the centralized vs. remote US scheduler deployment scenarios and the traverse paths of the REQ-GNT(MAP) messages.



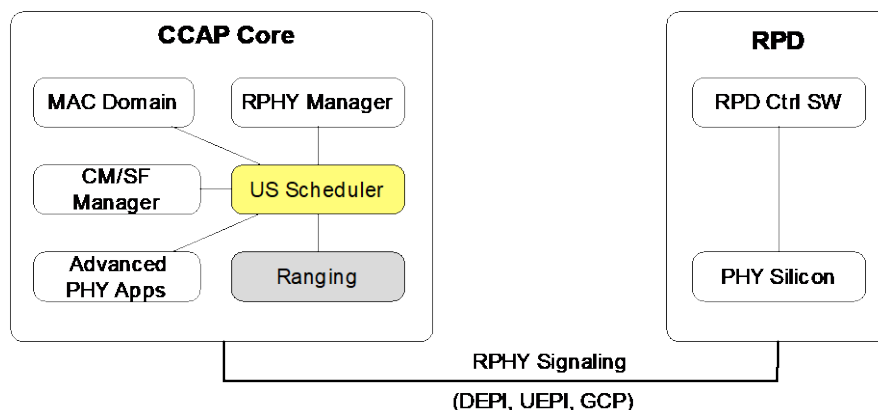
**Figure 1 – Centralized vs. Remote Upstream Scheduler Deployment Scenarios**

It should be noted that a centralized scheduler for Remote PHY will provide equivalent performance to a classic I-CCAP system as they both have the US scheduler in the same location. That location is the hub site that is 100 miles or less from the CM. The request and MAP messages are given high priority on the CIN so that any CIN queuing and latency will not add to the request grant delay.

Thus, the remote US scheduler is an operational option for when performance is desired that is better than a classic I-CCAP or if the CIN needs to be significantly longer than the 100 miles.

### 3. How to Spin off the Remote US Scheduler

In the R-PHY architecture today, the US scheduler is an internal component located at the CCAP Core, as shown in Figure 2.



**Figure 2 – Centralized Upstream Scheduler in R-PHY Today**

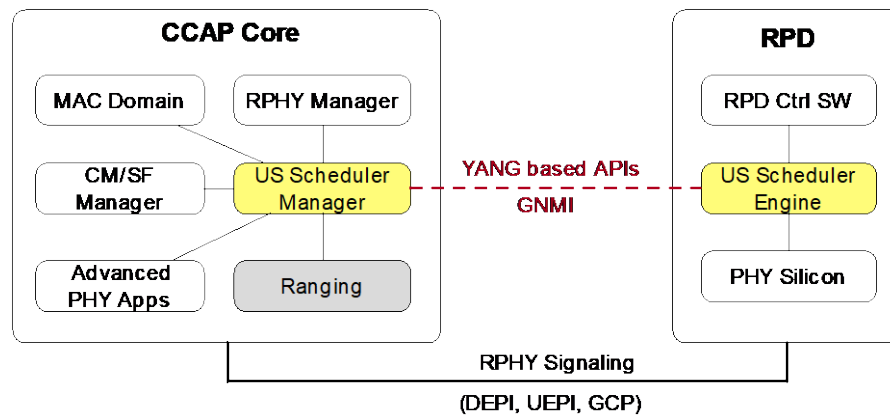
The US Scheduler interacts with various clients in the DOCSIS control plane using vendor specific interfaces, as listed in Table 1.

**Table 1 – US Scheduler Clients in DOCSIS Control Plane**

| Clients                   | US Scheduler Northbound Interface                                                                                                                                                                                            |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MAC Domain                | Notifies per MAC Domain upstream channel configurations including the corresponding primary capable downstream channels for carrying the MAP messages.                                                                       |
| Cable Modem (CM) Manager  | Notifies the CM operation state.                                                                                                                                                                                             |
| Service Flow (SF) Manager | Notifies the SF QoS Parameters, SF SID / SID cluster assignment.                                                                                                                                                             |
| Ranging                   | Requests for ranging transmission opportunities in MAP, including initial ranging, maintenance ranging and probing in OFDMA.                                                                                                 |
| Advanced PHY Applications | Requests for transmission opportunities for advanced PHY features, such as the data transmission opportunity for OFDMA US data profile (OUDP) test, or a silent transmission opportunity to capture per channel noise floor. |
| Remote PHY Manager        | Notifies the DEPI and UEPI session / pseudo wires to transport the scheduler signaling messages, such as MAP and bandwidth request pseudo wires.                                                                             |

Figure 3 shows a remote US scheduler architecture that separates out the real-time scheduling functions while keeping the northbound client interface intact. This is achieved by partitioning the US scheduler into a control component (US scheduler manager) in the CCAP core and a real-time component (US scheduler engine) in the RPD. The US scheduler manager interfaces the northbound control plane clients and manages the southbound US scheduler engines. The US scheduler engine fulfills the real-time scheduling services at the RPD.

In between the US scheduler manager and the US scheduler engine, we propose that a YANG data model-driven interface is used to formally describe the managed data source at the remote US scheduler engine. The US scheduler manager may simultaneously manage multiple US scheduler engines. To meet the timing and scaling requirements, a stream-lined YANG model-driven network protocol, such as gNMI, may be used to transport the API calls. gNMI stands for the gRPC Network Management Interface, and gRPC is known as Google remote procedure call.



**Figure 3 – Remote Upstream Scheduler Architecture**

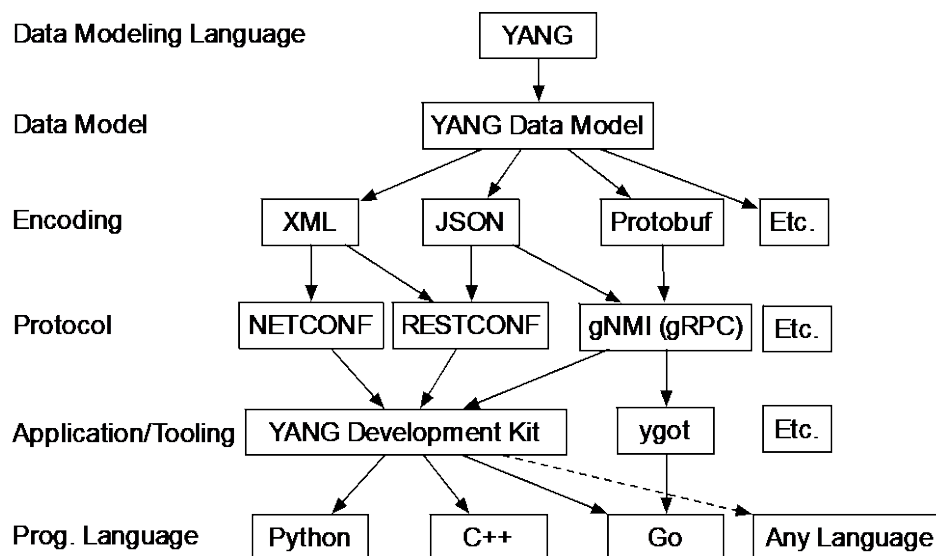
## 4. Why YANG Model-Driven API for the Remote US Scheduler

### 4.1. A Formal Definition of Service Contract

YANG is a full, formal contract language with rich syntax and semantics. A YANG model is the specification written in YANG about the managed data objects that a client can manipulate and observe using standardized operations. The YANG model-driven management is about building the YANG data and using the model via tooling to generate code to access the YANG objects on the managed devices.

As displayed in Figure 4, which covers the data model-driven management components, once the YANG models are specified and implemented, a particular encoding (XML, JSON, protobuf etc.) can be selected along with a particular remote procedure call (RPC) protocol (NETCONF, RESTCONF, or gNMI/gRPC) for transport. Based on the encoding and protocol selections, code or API can then be generated via proper applications / tooling (for example, in Python, C++, Go or any other language) [3].

To enable remote US scheduling, a formal service contract is needed between the US scheduler manager and the US scheduler engine. The YANG model-driven API is a natural choice to achieve this.



**Figure 4 – YANG Data Model-Driven Management Components**

## 4.2. Mature and Well Adopted

YANG was originally created at IETF in 2007 to describe standard data models for network automation via NETCONF. Since then, YANG data model-driven management has become a well-established trend in the network industry. YANG data models are produced by many parties including standard organizations, consortiums, forums and open source projects (such as OpenConfig and Broadband Forum), and network equipment vendors.

YANG data modeling has also been adopted in cable. There are several efforts involving YANG data modeling driven by CableLabs, including the ICCAP YANG modeling, as part of CCAP (Data-Over-Cable Service Interface Specifications) OSSI (Operations Support System Interface Specification) effort, the Flexible MAC Architecture (FMA) YANG modeling, and a recent initiative to build a common YANG ecosystem across different network elements (NEs), including CCAP Core, RPD, RMD and ICCAP.

## 4.3. Interoperation Made Easy

YANG solves the multi-vendor interoperability issue by using the data model as the core definition of the configuration and operational data and renders them to the management interface in a protocol specific construct. This approach is different from the traditional information model, which only models the managed objects at a conceptual level independent from the implementations. The YANG data model is intended for implementation. The YANG based API is bound to specific message encoding and data transport protocols that can be directly used by the client and server code. So as long as the underlying YANG modules are same, the client and server will have the same view of the API, avoiding the typical misinterpretation issues found at the interop.

## 4.4. Sufficiecnt Performance at Scale

The YANG based data model-driven management scheme supports a number of design options that can be used to achieve low-latency at scale. These include:

### Asynchronous and Parallel RPCs

By using asynchronous and parallel RPCs, a client can avoid the serialization delay and manage multiple devices or data objects at the same time. YANG natively supports asynchronous RPCs using the concept of intended configuration and applied configuration. This allows a client to acquire the RPC result asynchronously by either polling the applied configuration or subscribing to the applied configuration update, without being blocked by the RPC results. Additionally, the client and server can use the YANG data tree to identify the interdependencies among the data objects and enable parallel processing of the data objects are isolated from each other.

### **Long-lived RPC sessions**

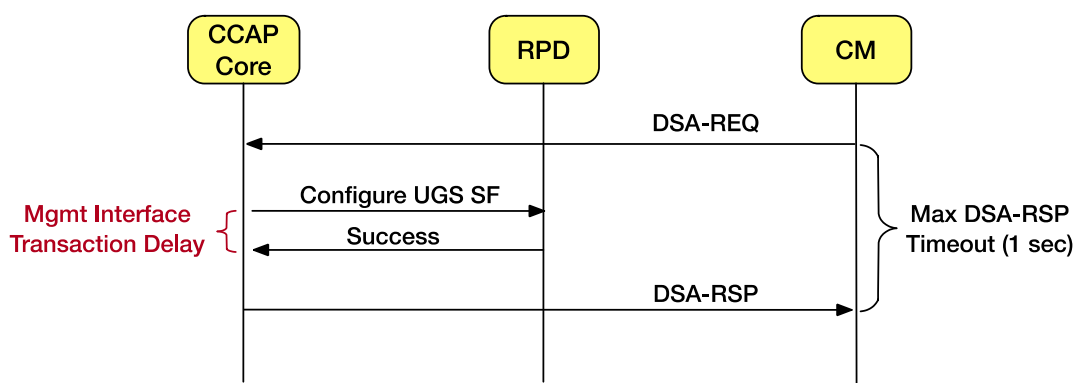
Certain YANG based protocols, such as NETCONF and gNMI (gRPC), support long-lived RPC sessions to minimize session setup overhead. gNMI (gRPC), which uses HTTP/2, further allows multiple RPCs to be multiplexed onto one TCP connection. By doing so, it supports parallel RPCs without increasing the number of TCP connections.

### **Binary Encoding**

gNMI (gRPC) supports Protobuf, an efficient binary message format, that improves message encoding/decoding processing time on both client and server. Based on Google's gRPC performance benchmark test, 700k unary RPC calls per second is achievable [5].

As to the remote US scheduler interface, the most stringent timing constraints come from two basic requirements. These are the dynamic service flow setup for voice, and the need for continuous ranging.

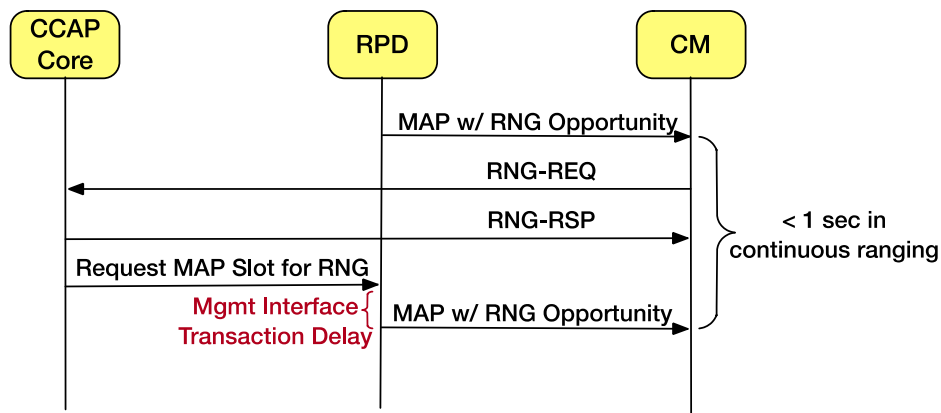
Let's first examine the dynamic service flow setup case as shown in Figure 5. After receiving the DSA-REQ for a voice call, the scheduler manager in the CCAP core needs to check with the US scheduler engine in the RPD to make sure the required quality path can be established. In this case, the YANG based API transaction delay contributes to the overall service flow setup time, which needs to be one second or less to meet the post-dial delay requirement [4].



**Figure 5 – Adding a UGS SF in a CM Initiated DSA Transaction**

Now let's see about ranging as shown in Figure 6. After receiving the RNG-REQ from the CM, the ranging module in the CCAP core may decide to continuously range the CM by requesting a subsequent ranging opportunity in less than a second if it detects the CM deviates from the optimum timing/power/frequency settings.





**Figure 6 – Adding a RNG-REQ Opportunity in MAP**

For both cases, the overall DOCSIS transaction delay is capped at one second or less. To meet this requirement, we can assume tighter budget of 100-millisecond delay budget for the YANG based remote US scheduler management interface, taking into consideration of the typical RPD CPU speed, CIN delay, and RPC transaction time.

#### 4.5. Easy to Maintain

One advantage for using YANG based management interface is the rich tooling and applications for processing YANG modules and their metadata. YANG applications validate YANG modules, generate APIs and provide language binding. YANG metadata allows the client to pre-validate the instance data, confirm the module support in the network elements, and check for potential non-backward-compatible changes that could have been introduced between versions. YANG also has a huge library contributed by various organizations that promotes reusability.

YANG model driven telemetry (MDT) provides a new way to maintain service quality by streaming data continuously from the managed devices using a push model to give client near real-time access to the operational statistics.

### 5. Remote Upstream Scheduler Behavior Model

The goal is to place a remote upstream scheduling engine into an RPD that works in tandem with a remote upstream scheduling manager in the CMTS Core. Aside from the specific functionality of this system, there are some high-level objectives that have to be met to turn this into a product. Some of these at least are:

- It has to work
- It has to scale
- It has to interoperate
- It has to be maintainable

The following discussion is aimed at hitting all of these goals.

One option would just to let different manufacturers load their US schedulers from their cloud CMTS into someone else's RPD. In theory that could work, and may well be a product option. But then there is software from two manufacturers in one platform which is challenging from a built and test viewpoint.

Instead, this paper proposes building an US scheduler from one manufacturer that will interoperate with another manufacturer.

This paper proposes two fundamental principles for constructing a remote scheduler.

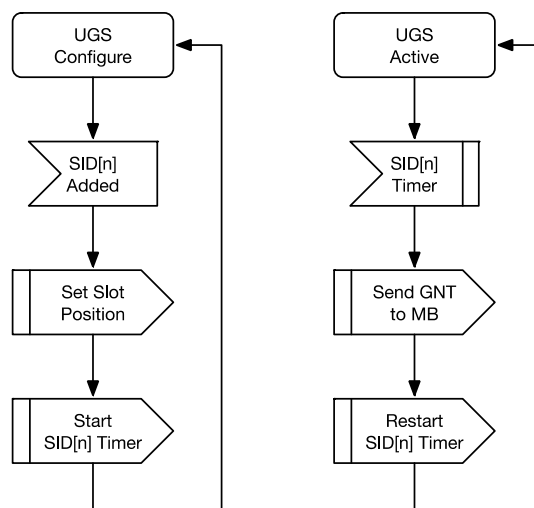
1. Create a well-defined behavior model
2. Use a well-defined API on that behavior model

In the construction of this behavior model and API, a very conscience decision was made to take the larger scheduler and break it down into smaller independent schedulers that each had very specific functionality. There would then be a rule set on how the different schedulers interacted with each other. Finally, there would be another behavior model for the MAP Builder.

It's time for an example. Let's look at the behavior model for the unsolicited grant scheduler (UGS).

#### UGS Service Flow

- Service Flow ID (SFID)
- MAC domain name
- MAC domain upstream channel
- Service Identifier (SID)
- UGS QoS Parameters
  - request-policy
  - unsolicited grant size (bytes)
  - nominal grant interval (ms)
  - tolerated grant jitter (ms)
  - grants per Interval
- Grant statistics
  - cumulative grant size
  - grant rate



**Figure 7 – UGS Behavior Model**

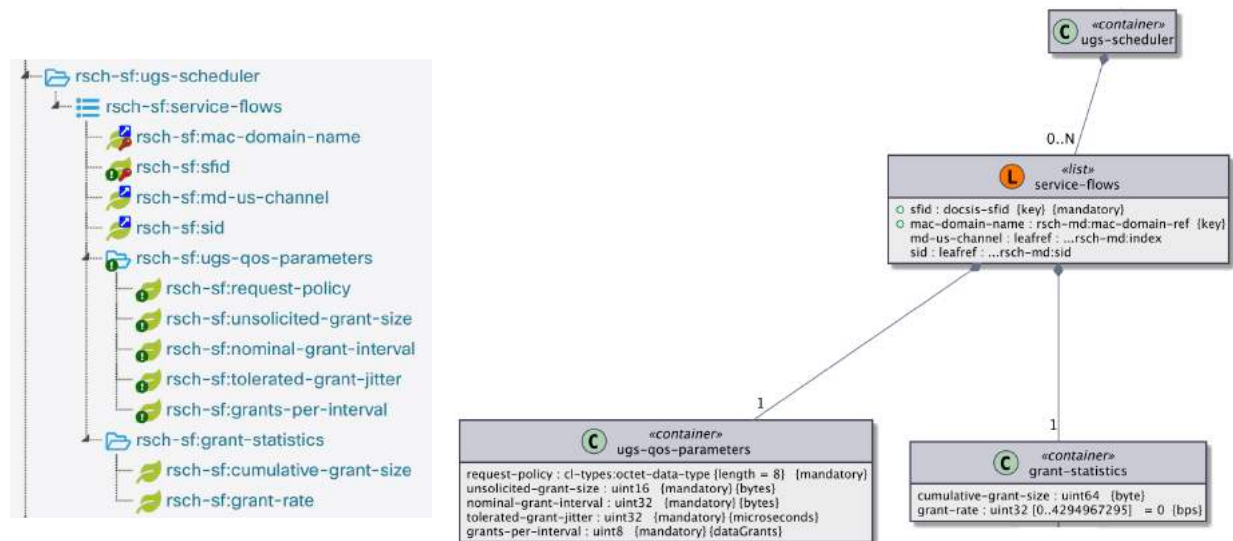
A behavior model and a basic API is shown in Figure 7. This API will ultimately be represented in YANG but is shown here in variable form.

The behavior model has two states, configure and active. The configure state is for adding and removing service flows. The active state is run-time that generates the grants when needed. The model is simple. A service flow/SID is added into a scheduling wheel. For example, if a grant for voice-over-IP needs to be sent every 20 ms and a MAP message is every 2 ms, then there are 10 entries in the scheduling wheel. At run time, an event timer expires and a grant is scheduled.

Now let's look at the API. There are three basic elements. First, there is the logical associations. This is done with a service flow ID, a mac domain name, an upstream channel identifier and a SID. Second, is what action to take. This is described by the QoS parameters. The action is then with a request policy of UGS, to prohibit any request opportunities, instead, directly provide an unsolicited grant size (say 100 bytes) every nominal grant interval (say 20 ms) with a tolerated jitter (say 2 ms) with a grants per interval (say 1). Third, it is important to always measure what happened. In this example, the grant statistics include the cumulative grant size and the grant rate.

So, the API connects the two sides together, an action is requested, and a measurement is made to see how the action played out. That API describes a behavior model. That's it. This methodology can be repeated for the other schedulers and for the MAP builder.

In Figure 8, we see the actual implementation of the final model in the YANG tree and UML.



**Figure 8 – UGS Diagram and UML Model**

To go back to our original objectives, the model has to work. That is achieved through simplicity and modularization of the model. The model has to scale. That is achieved by setting realistic performance and scaling goals. The model has to interoperate. That is achieved by using YANG as the API and well-defined behavior models. And the model has to be maintainable. That is achieved again with YANG and the MDT features that allow monitoring and measuring of what has happened.

## 6. YANG Model for Remote Upstream Scheduler

The remote US Scheduler YANG model includes a base module that describes the remote US scheduler framework and several submodules that specifies the component level attributes. The overall remote US scheduler YANG model is intended for the following uses cases:

- Configure the remote US scheduler, e.g.,
  - Configure US SF scheduling type and QoS parameters
  - Configure US MAP attributes
- Invoke US scheduling actions, e.g.,
  - Request an US ranging opportunities in MAP
- Collect the remote US scheduling telemetry data, e.g.,
  - Acquire per SF requesting rate and granting rate

## 6.1. Base remote-us-scheduler YANG Module

The base module, titled “remote-us-scheduler” defines the remote US scheduler structure that holds together three functional submodules, namely us-scheduler-domain, us-qos-scheduler and map-builder, as shown in Figure 9. Each of the submodules can be defined separately and augmented into the base module. This arrangement decomposes a complex scheduler into smaller and simpler submodules that can be independently modeled and managed.

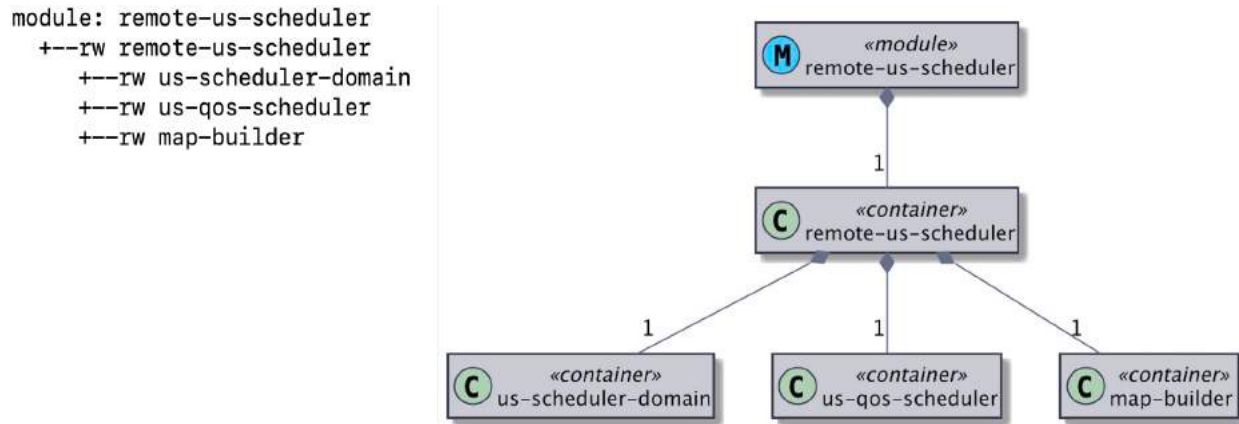


Figure 9 – remote-us-scheduler base module tree diagram and UML

## 6.2. us-scheduler-domain YANG Module

This module defines the MAC domains served by the remote US scheduler. Each MAC domain contains the parameters required for US scheduling, including common QoS policies, US channels and SID assignments and the list of primary capable DS channels for carrying MAP messages.

Figure 10 shows the top-levels of the us-scheduler-domain YANG module tree diagram, omitting the lower level leaf entries for simplicity. The us-scheduler-domain module is augmented to the us-scheduler-domain container in the remote-us-scheduler base module.

```

module: us-scheduler-domain
 augment /rsch:remote-us-scheduler/rsch:us-scheduler-domain:
 +--rw us-scheduler-domain* [mac-domain-name]
 +--rw mac-domain-name string
 +--rw md-us-channel* [index]
 | +--rw index uint8
 | +--rw us-rf-chan-cfg
 | | ...
 | +--rw ds-rf-chan-cfg* [port-number channel-index]
 | | ...
 | +--ro md-us-channel-state
 | | ...
 | +--rw sid-cfg* [sid]
 | | ...
 +--rw md-us-qos-policy* [type]
 +--rw type identityref
 +--rw (parameters)?

```

Figure 10 – us-scheduler-domain module top-three level tree diagram

### 6.3. us-qos-scheduler YANG Module

This module defines the US QoS scheduler of different scheduling types, including best effort, rtPS, nrtPS, UGS, PGS and aggregated service flows (ASF). Each scheduler type represents a specific scheduler behavior as discussed in Section 5.

Figure 11 shows the top-level tree diagram of the us-qos-scheduler YANG module, omitting the lower level leaf entries for simplicity. Figure 12 shows a detailed tree-diagram for the best-effort-scheduler. The us-qos-scheduler module is augmented to the us-qos-scheduler container in the remote-us-scheduler base module.

```
module: us-qos-scheduler
 augment /rsch:remote-us-scheduler/rsch:us-qos-scheduler:
 +--rw best-effort-scheduler
 | ...
 +--rw rtps-scheduler
 | ...
 +--rw nrtps-scheduler
 | ...
 +--rw ugs-scheduler
 | ...
 +--rw ugs-ad-scheduler
 | ...
 +--rw pgs-scheduler
 | ...
 +--rw lld-asf-scheduler
 | ...
 +--rw dhqos-asf-scheduler
 | ...
```

**Figure 11 – us-qos-scheduler module top level tree diagram**

```

module: us-qos-scheduler
augment /rsch:remote-us-scheduler/rsch:us-qos-scheduler:
 +--rw best-effort-scheduler
 | +--rw service-flows* [mac-domain-name sfid]
 | | +--rw sfid docsis-sfid
 | | +--rw mtc-mode-enabled? boolean
 | | +--rw mac-domain-name rsch-md:mac-domain-ref
 | | +--rw (sid-channel-assignment)?
 | | | +--:(non-mtc)
 | | | | ...
 | | | +--:(mtc)
 | | | | ...
 | +--rw best-effort-qos-parameters
 | | +--rw request-policy? cl-types:octet-data-type
 | | +--rw priority? uint8
 | | +--rw data-rate-unit-setting? cl-fma-qos:data-rate-type
 | | +--rw max-traffic-rate? uint32
 | | +--rw max-traffic-burst? uint32
 | | +--rw min-reserved-rate? uint32
 | | +--rw min-reserved-packet? uint16
 | | +--rw max-concatenated-burst? uint16
 | | +--rw peak-traffic-rate? uint32
 | +--ro request-grant-statistics
 | | +--ro cumulative-request-size? uint64
 | | +--ro request-rate? uint32
 | | +--ro cumulative-grant-size? uint64
 | | +--ro grant-rate? uint32

```

**Figure 12 – best-effort-scheduler tree diagram**

## 6.1. Map-builder YANG Module

This module defines the MAP builder for the remote US scheduling. It contains MAP builder configuration and operational data, as well as RPC calls as shown in Figure 13.

The MAP building configuration parameters reflect channel level MAP attributes such as ranging and data back off window, broadcast IM and idle slot settings. The MAP builder RPCs are used to invoke dynamic MAP actions, such as building a station maintenance (SM) ranging opportunity or request for an OFDMA upstream data profile (OUDP) test slot in MAP. The input and output parameters allow the client to specify the action requirement and retrieve data coming out of the action. Unsolicited notifications can also be added to report the MAP builder operational status. Alternatively, the client can subscribe the operational data using telemetry.

```

module: map-builder
 augment /rsch:remote-us-scheduler/rsch:map-builder:
 +--rw map-builder* [mac-domain-name md-us-channel]
 +--rw mac-domain-name rsch-md:mac-domain-ref
 +--rw md-us-channel -> /rsch:remote-us-schedule:
 +--rw map-channel-enable enumeration
 +--rw ranging-backoff-start uint8
 +--rw ranging-backoff-end uint8
 +--rw transmit-backoff-start uint8
 +--rw transmit-backoff-end uint8
 +--rw broadcast-im-cfg
 | ...
 +--rw padding-slot-cfg
 | ...
 +--ro map-channel-state
 | ...

rpcs:
 +--x udp-test-slot-request
 | +--w input
 | | ...
 | +--ro output
 | | ...
 +--x unicast-im-request
 | +--w input
 | | ...
 | +--ro output
 | | ...
 +--x sm-request
 | +--w input
 | | ...
 | +--ro output
 | | ...
 +--x probe-request
 | +--w input
 | | ...
 | +--ro output
 | | ...
 +--x idle-slot-request
 | +--w input
 | | ...
 | +--ro output
 | | ...

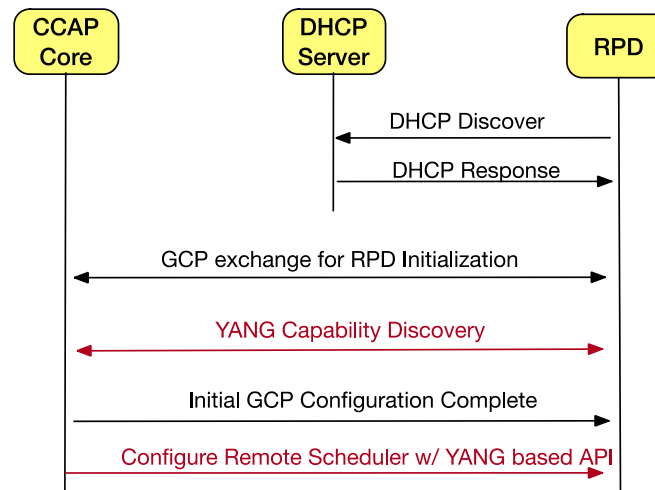
```

**Figure 13 – map-builder top level tree diagram**

## 7. Coexistence with Existing RPD Management Interface

In R-PHY today, the management interface runs over the Generic Control Protocol (GCP) using the Type-Length-Value (TLV) tuples to carry the signaling messages defined in the R-PHY specification [6]. Since the remote US scheduler is a new functional component, running GCP as it is today does not prevent using the YANG based API for the remote US scheduler, as the two management interfaces do not manipulate the same data objects.

Figure 14 shows the initialization sequence with the YANG based management interface coexisting with the legacy GCP based RPD management interface. Once the RPD is assigned with an IP address and paired with the CCAP core, the core can initiate the YANG capability discovery process. For an RPD that declares the remote US scheduler YANG support, the core can program the scheduler engine once the RPD is initialized. If there are any configuration dependencies on the GCP managed objects, for example the RF configuration, the US scheduler engine at the RPD can hold off the intended configuration and only apply it when the GCP dependencies are resolved.



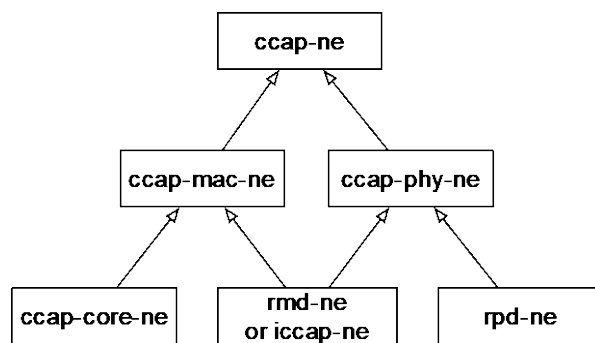
**Figure 14 – Coexistence with Existing RPD Management Interface**

With this hybrid management model, the remote US scheduler can be deployed on existing RPDs without waiting for the full-blown RPHY YANGification.

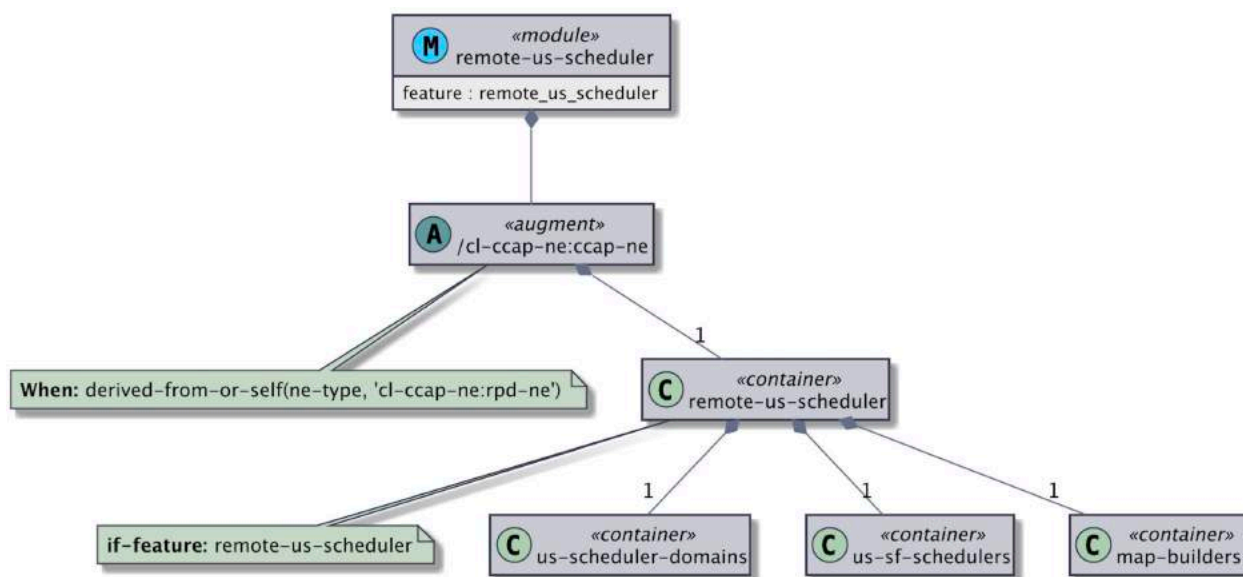
## 8. Remote US Scheduler in DOCSIS YANG Ecosystem

In the DOCSIS YANG ecosystem, different CCAP network elements (NEs) can be organized into classification hierarchies as shown in Figure 15, where the rpd-ne (representing RPD NE) is one type of the CCAP NE that inherits the ccap-phy-ne attributes. The remote US scheduler YANG module can be incorporated into the CCAP NE hierarchy by augmenting the rpd-ne base module as shown in Figure 16. The “when” and “if-feature” tags allow the remote US scheduler module to be only present if the CCAP-NE is the RPD type and supports remote US scheduling.





**Figure 15 – CCAP Network Element Types and the Subclassing Hierarchy**



**Figure 16 – Remote US Scheduler Module in DOCSIS YANG Ecosystem**

## 9. Conclusion

The remote US scheduler is an operational option for R-PHY when the desired latency performance needs to be better than a classic I-CCAP or if the CIN needs to be significantly longer than the 100 miles. The remote US scheduler design has two functional parts, the US scheduler manager that lives in the CCAP core and the US scheduler engine that lives in the RPD. The US scheduler manager northbound is internal to the CCAP core, while the southbound is a data model driven interface that allows the US scheduler manager to communicate with one or more US scheduler engines to fulfill the US scheduling services.

This paper proposes to use a YANG data model-driven interface between the US scheduler manager and the US scheduler engines. The YANG data-model driven management techniques are mature and well adopted in the network industry, with rich tooling and applications that can auto-generate code from the data model; hence facilitating multi-vendor interoperability. The YANG data model can be paired with high performance message encoding and RPC options (such as protobuf and gRPC) to meet the timing and scaling requirement for managing remote US schedulers.

To construct the remote US scheduler YANG data model, this paper presents a modeling method that involves decomposing a complex service function into smaller independently manageable subservices, creating a well-defined behavior model for each subservice, and defining the API that completely and precisely characterizes the behavior.

The YANG model-driven management interface for remote US scheduler can co-exist with existing GCP based RPD management interface, reducing time-to-market for the remote US scheduler feature while R-PHY is transitioning to YANG data-model driven management. The remote US schedule YANG model is an integral part of the DOCSIS YANG ecosystem, it is planned to be contributed to CableLabs as a feature augmented to the R-PHY YANG model.

## Abbreviations

|        |                                               |
|--------|-----------------------------------------------|
| API    | application programming interface             |
| CM     | cable modem                                   |
| CCAP   | converged cable access platform               |
| CIN    | converged interconnect network                |
| DS     | downstream                                    |
| DSA    | dynamic service addition                      |
| FMA    | flexible MAC architecture                     |
| GCP    | generic control protocol                      |
| gNMI   | gRPC network management interface             |
| GNT    | DOCSIS bandwidth grant                        |
| gRPC   | Google RPC                                    |
| MAP    | DOCSIS bandwidth request                      |
| NE     | network element                               |
| RPC    | remote Procedure Call                         |
| RPD    | remote PHY device                             |
| RPD-NE | RPD Network Element                           |
| YANG   | yet another next gen (data modeling language) |

## Bibliography & References

- [1] Tong Liu, John Chapman, “R-PHY with Remote Upstream Scheduler”, *2019 SCTE Expo Technical Forum Proceedings*, Oct, 2019.
- [2] “CM-SP-MULPIv3.1-I18-190422: MAC and Upper Layer Protocols Interface Specification”, CableLabs, 2019
- [3] Joe Clarke, Jan Lindblad, Benoit Claise: *Network Programmability with YANG* Addison-Wesley Professional Book
- [4] ITU-T: Series J: Cable Networks and Transmission of Television, Sound Programm and Other Multimedia Signals
- [5] <https://grpc.io/docs/guides/benchmarking/>
- [6] “CM-SP-R-PHY-I14-200323: DOCSIS Remote PHY Specification”, CableLabs, 2020

# **Training Machines to Learn From Signal Meter Readings**

## **A Case Study from Comcast**

A Technical Paper prepared for SCTE•ISBE by

**Gary Ventriglia**

Sr. Principal Engineer  
Comcast Corporation  
1800 Arch St, Philadelphia PA 19103  
484-846-0467  
Gary\_Ventriglia@comcast.com

**Jack Birnbaum**

Vice President, Customer Experience  
Comcast Corporation  
1800 Arch St, Philadelphia PA 19103  
215-286-8057  
Jack\_Birnbaum@cable.comcast.com

**Robert Gonsalves**

Sr Director, Product Dev & Engineering  
Comcast Corporation  
1401 Wynkoop Suite 300, Denver CO 80202  
720-512-3642  
Robert\_Gonsalves@cable.comcast.com

**Anastasia Vishnyakova**

Engineer 3, Machine Learning  
Comcast Corporation  
1800 Arch Street, Philadelphia PA 19103  
267-260-3277  
Anastasia\_Vishnyakova@comcast.com

**Michael Kreisel**

Principal Researcher, Machine Learning  
Comcast Corporation  
1110 Vermont Ave NW, Washington, DC 20005  
202-524-5068  
Michael\_Kreisel@comcast.com

**Larry Wolcott**  
Comcast Fellow of Engineering  
Comcast Corporation  
1401 Wynkoop Suite 300, Denver CO 80202  
720-512-3643  
Larry\_Wolcott@cable.comcast.com

# Table of Contents

| Title                                                     | Page Number |
|-----------------------------------------------------------|-------------|
| 1. Introduction.....                                      | 5           |
| 2. Acknowledgements.....                                  | 5           |
| 3. A Brief History of RF Troubleshooting.....             | 5           |
| 4. Fault Isolation on RF Transmission Lines.....          | 6           |
| 4.1. The Troubleshooting Process.....                     | 10          |
| 4.2. Repair Feedback.....                                 | 11          |
| 4.2.1. Technician Performance Metrics.....                | 11          |
| 4.2.2. Customer Relationships.....                        | 11          |
| 4.2.3. The Expert Technician / Opportunistic Repairs..... | 12          |
| 4.2.4. Report Burden and Selection Bias.....              | 12          |
| 4.2.5. Intermittent Problems and Upstream Noise.....      | 12          |
| 4.2.6. Insufficient Tools or Training.....                | 13          |
| 5. Cloud Connected Technicians and Equipment.....         | 13          |
| 5.1. A New Troubleshooting Process.....                   | 13          |
| 5.2. Process Compliance.....                              | 14          |
| 6. Machine Learning and Artificial Intelligence.....      | 15          |
| 6.1. Features and Labels.....                             | 15          |
| 6.1.1. Machine Data vs Human Data.....                    | 16          |
| 6.1.2. Digital Interactions With our Customers.....       | 17          |
| 6.2. The Model.....                                       | 17          |
| 6.2.1. Outside of Home Applied AI Model.....              | 17          |
| 6.2.2. Using Weak Supervision to Improve the Model.....   | 20          |
| 7. Conclusion.....                                        | 23          |
| Abbreviations.....                                        | 24          |
| Bibliography & References.....                            | 25          |

## List of Figures

| Title                                                         | Page Number |
|---------------------------------------------------------------|-------------|
| Figure 1 – A Typical RF Troubleshooting Process.....          | 6           |
| Figure 2 – Common Problem Signature, Multiple Devices.....    | 7           |
| Figure 3 – Single Problem Signature, Multiple Locations.....  | 8           |
| Figure 4 – Problem Signature Compared to Neighbors.....       | 8           |
| Figure 5 – Localization Based on Signature Inference.....     | 9           |
| Figure 6 – Upstream Noise Impact on Entire Service Group..... | 10          |
| Figure 7 – New Troubleshooting Process.....                   | 14          |
| Figure 8 – Example of Cable-Oriented Features and Labels..... | 16          |
| Figure 9 – Node Topology Schema.....                          | 19          |
| Figure 10 – Model Performance, Precision and Recall.....      | 20          |
| Figure 11 — Model Feature Importance.....                     | 20          |
| Figure 12 – Weak Supervision Workflow.....                    | 21          |
| Figure 13 – ML/AI Enhanced Process and Feedback Loop.....     | 24          |

## List of Tables

| <b>Title</b>                                                      | <b>Page Number</b> |
|-------------------------------------------------------------------|--------------------|
| Table 1 – Model Sources Summary .....                             | 18                 |
| Table 2 – Labeling Functions with Highest Empirical Accuracy..... | 22                 |
| Table 3 – Model Performance Comparison .....                      | 23                 |

## 1. Introduction

When informed by vast amounts of network performance information, identifying radio frequency (RF) problems with the Data-Over-Cable Service Interface Specifications (DOCSIS) isn't that hard. However, determining if the problems are inside or outside the home can be difficult. This is a decades-old problem, with hopes often pinned on the elusive promise of artificial intelligence (AI) or machine learning (ML) to help. The challenge that many data scientists will tell you is that having good training data is critical. The lack of a reliable feedback loop to establish cause-and-effect often results in poorly trained machines.

A significant amount of time and resources has been poured into remote diagnostic tools to identify plant problems. Those tools historically have been segmented, specialized, and tuned to evaluate singular aspects of the RF health – for example, receive modulation error ratio (RxMER) and forward error correction (FEC). Once a problem is identified, determining if its source is in a customer's home, drop or tap has historically been left to technicians, to provide feedback about what they found. The feedback mechanisms typically involve selecting a code or result and updating the work order when it's complete.

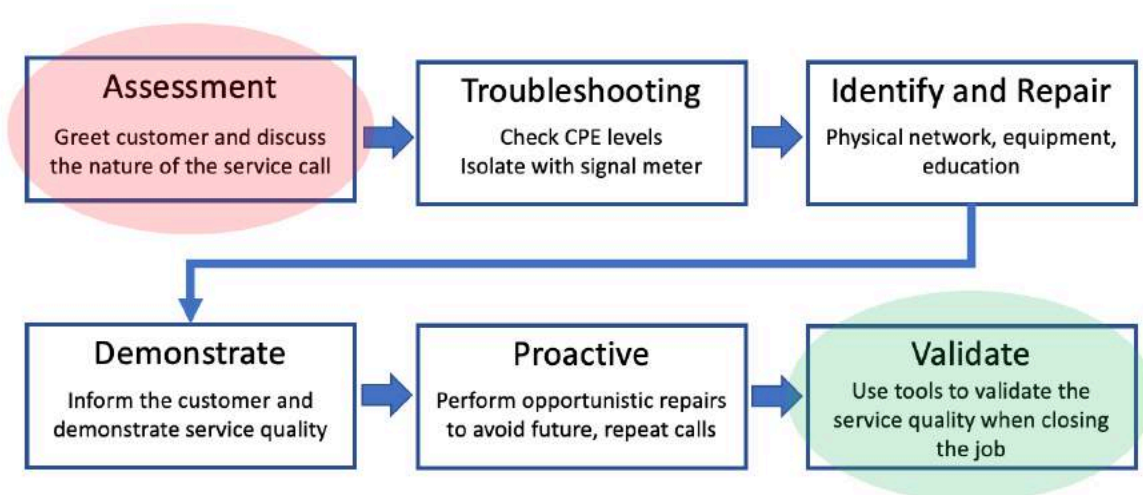
With the COVID-19 pandemic starting mid-March 2020, the rapid development of an “outside network check” provided an opportunity to gather better features and labels. With a renewed desire to keep technicians and customers isolated, the team is exploring new ML/AI models. These new models are trained to use cloud-based RF measurements. These measurements include remote telemetry from DOCSIS devices, and other equipment logged by collection systems. Another set of measurements is taken at the tap and ground block, finally offering a way to segment the network and train the machines differently. The authors review the outcome of this fascinating exercise currently under way, as this paper is being written, in the summer of 2020.

## 2. Acknowledgements

Much of the material within this document was obtained from interviews and written contributions of many Comcast leaders and other industry experts. The authors would like to thank David Monnerat, Gary Schwin, Justin Menard, Shawn Hughes, Patrick McDonald, Miles Pellegrini, Marty Marcinczyk, Jan Neuman, Rama Mahajanam, and Fan Liu from Comcast. You all collectively represent a large cross-section of unique perspectives and expertise. We also recognize Brady Volpe of the Volpe Firm and his award-winning products and services. Thank you, Brady, for your ongoing industry leadership and inspiration. Finally, we express gratitude to our distinguished colleagues Seamus Gallagher, Dwain Kelly and Paul Kelly from Liberty Global. Your global contributions and experience are invaluable to our industry.

## 3. A Brief History of RF Troubleshooting

For well over 40 years, cable field technicians have relied on signal level meters (SLMs) to take measurements that help determine a proper course of repair. Once alerted and dispatched to a problem, our technicians have been trained to use a divide-and-conquer approach for troubleshooting, and ultimately repairing the issue. These meters would usually be used to take measurements of RF signals to-and-from the customer at different locations on the network. They might start inside the customer location, to validate the service at the location of the customer's equipment, then work their way towards the network to determine where the problem begins or ends. There are other processes where technicians might start at the network tap and work their way towards the customer equipment. Of course, these are basic processes which traditionally made it very difficult to verify if, or how they were followed.



**Figure 1 – A Typical RF Troubleshooting Process**

Courtesy of Doug Kelly, Virgin Media, Ireland

Meanwhile, the cable industry has been investing in remote performance telemetry for a long time, beginning with the DOCSIS 1.0 release in 1997. For nearly 23 years, the DOCSIS specification has been evolving and growing, providing performance metrics for virtually every aspect of our cable networking protocols.

Since DOCSIS 3.0 was launched in 2006, the cable industry has been developing a proactive network maintenance (PNM) specification, which significantly improves the range and depth of available RF diagnostic capabilities. In the hands of a highly skilled technician, these tools can be invaluable and empowering. However, in some cases, they require certain skills and experience, in order to properly interpret and enact a correct repair. When a technician lacks the training or experience to properly use or interpret the tools, they often forego using them and rely a limited repertoire for repair, such as swapping equipment and resetting devices. Our technical workforce is made up of skills on both sides of this continuum and everywhere in between.

As the remote visibility of the customer equipment improves, so do the tools to detect problems, present data and dispatch technicians. Quoting Brady Volpe, owner of The Volpe Firm Inc., “In the mantra of PNM, we are often able to find and repair problems before they impact the performance of the service perceived by our customers.” While that’s an attractive statement to make, it can be difficult to determine exactly which problem is causing the trouble being experienced by the customer. We may be able to proactively detect RF problems before the customer is impacted, but they could be calling due to a completely unrelated problem. In many cases, our tools currently lack the ability to accurately establish cause-and-effect with the customer experience.

## 4. Fault Isolation on RF Transmission Lines

To people familiar with customer support and troubleshooting, there are a few relatively simple processes that tend to work universally. Sometimes, problem isolation can follow a basic divide-and-conquer approach, depending on the type of problem. However, RF transmission troubleshooting has special conditions and circumstances that can be problematic for a simplistic troubleshooting process.



It may be helpful to remember that our coaxial cables are essentially just a large, shielded transmission line. Our shielded transmission lines can cover long distances, spanning many splits and taps with little directional isolation, which can befuddle the divide-and-conquer approach. Fortunately, cable has some physical characteristics that help, but do not completely solve these isolation problems. It can be generalized that many problems are difficult or impossible to localize, inside or outside of the customer premises.

The most common and reliable method for automatically localizing problems within a location is the presence of multiple equipment. This allows the remote monitoring systems (and technicians) to compare the signals of multiple sensors within the location. By a process of comparison, all devices within the home sharing a common problem increases the likeliness that the problem is outside (Figure 2). Conversely, if a single device detects a problem while the others do not, it can be concluded that the problem is inside. In recent years, a boom in DOCSIS deployments has made this technique very useful. However, the latest industry trend is moving towards a single DOCSIS point-of-entry gateway, then relying on Wi-Fi to distribute content within the premises. Over time, this will diminish the value of this localization technique.




**Figure 2 – Common Problem Signature, Multiple Devices**

In a similar process of elimination, when a single problem signature is detected having multiple sensors available (Figure 3), it's a reasonable assertion that the problem is within the location. These types of issues are typical of wiring problems such as loose or damaged connectors, incorrect fittings on wall plates, damaged cables, and splitters that are installed backwards, to name a few. In some cases, the problems can be addressed by the customer without requiring a technician. Some operators use this technique to identify possible loose connectors and inform the customer, instructing them to tighten the connector. These types of problems may not cause a noticeable issue for the customer, but tightening loose connectors is one of the easiest and best ways to improve the overall reliability of a coaxial network.



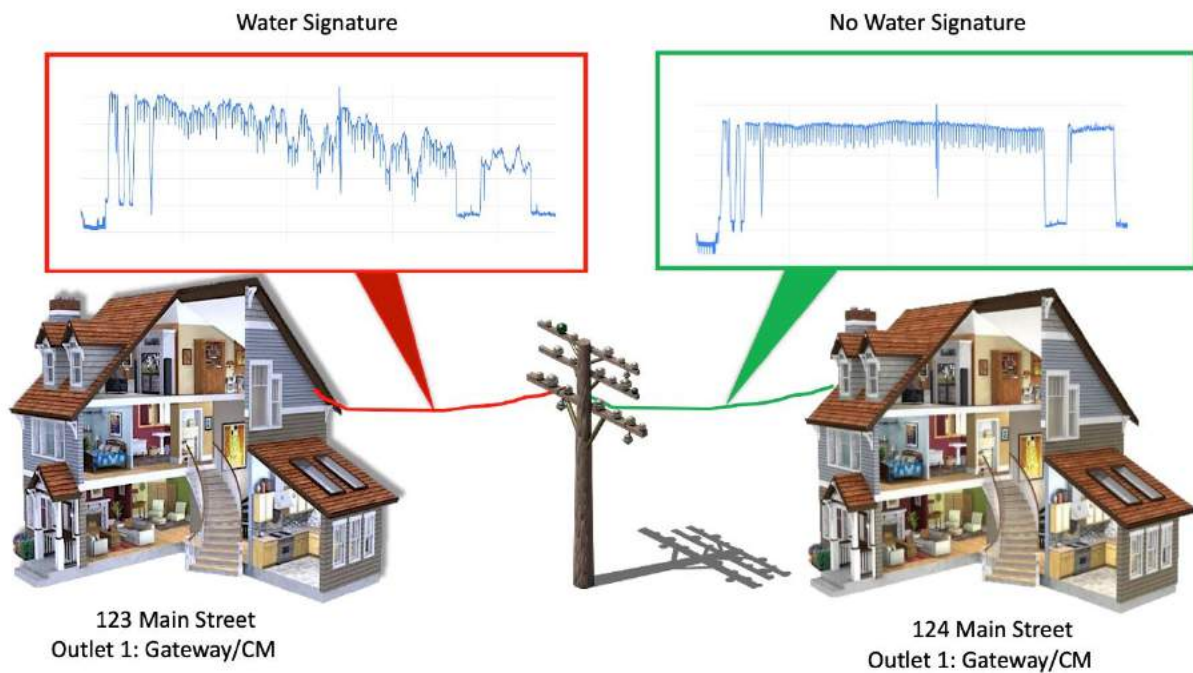
**Figure 3 – Single Problem Signature, Multiple Locations**

The technique for comparison-based localization improves as more data is introduced, including sensor information from neighboring equipment connected to the same physical network tap (see Figure 4). When common problems are detected that have a shared network element, the shared network element is usually the cause for impairment. These types of problems are commonly found at the tap, and include damaged or corroded tap plates, incorrect termination, cut drops, incorrect pin length, loose pin seizures and more.



**Figure 4 – Problem Signature Compared to Neighbors**

Another useful method for localizing problems uses casual observation and makes inferences based on common signatures. For example, there are certain impairments which are well-known to be outside vs. inside. In the case of water damage, water enters a coaxial cable and creates a distinctive signature that is easily identified (Figure 5). These are nearly always outside, caused by environmental influences such as rain or sprinkler systems. There are possible examples where water can have entered coaxial inside the home, but those are minimal and unlikely. Other examples of inference-based signatures are the presence of filters, which are usually installed at the ground block or tap; old satellite splitters which can produce unique standing waves; and spectral roll-off from old passive equipment, to name a few. These types of signatures are automatically detected with specially designed spectral impairment detection (SID) software libraries available from CableLabs.

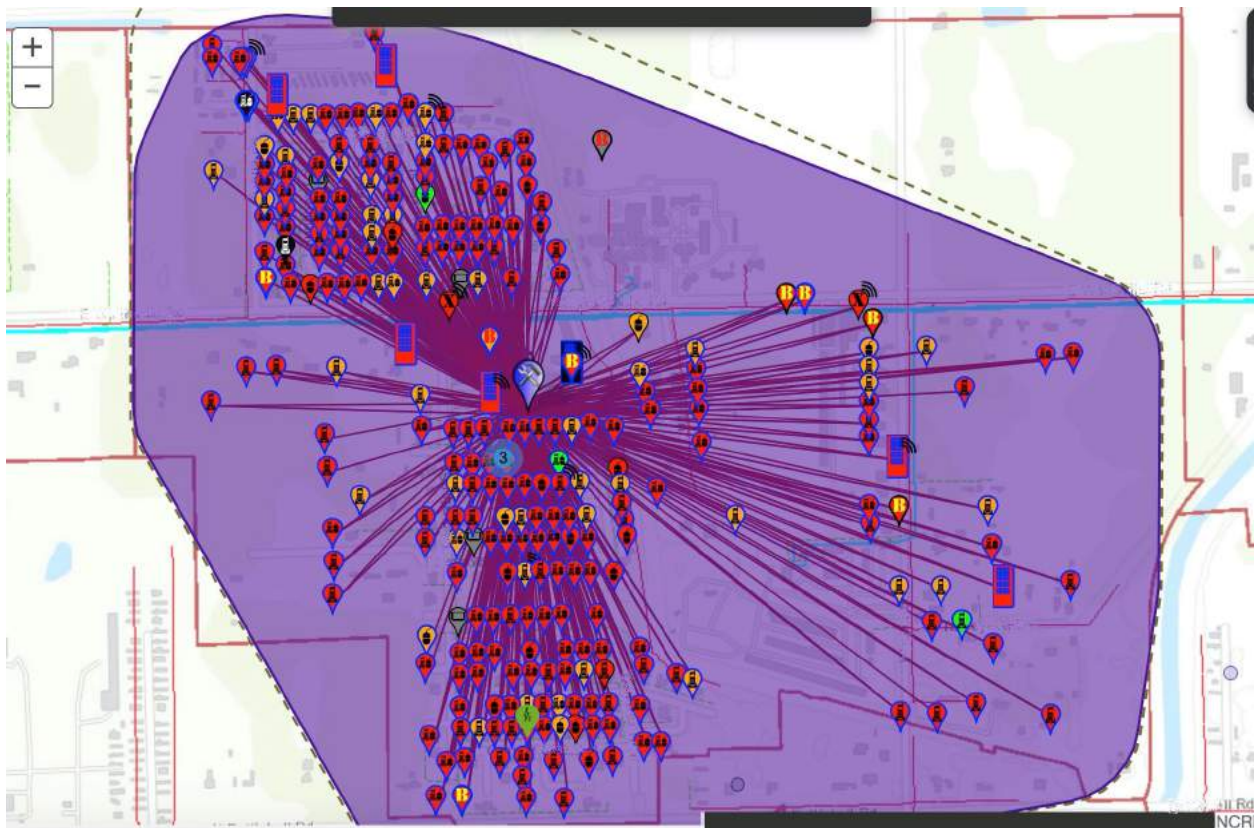


**Figure 5 – Localization Based on Signature Inference**

The techniques described above work on a subset of the common types of problems that are detectable with our DOCSIS PNM tools. They are especially useful for downstream RF impairments and anything that causes an impedance mismatch, any physical damage on the cable system. As previously discussed, these techniques do not work for certain types of problems. Unfortunately, some of our most impactful and common problems are invisible to this type of isolation technique.

Upstream noise is the most notable exception, including any other spurious types of interference which may be intermittent. The primary reason for upstream noise eluding these techniques is known as the upstream funnel effect. The upstream portion of the RF spectrum is coupled in a manner that allows all of the signals to travel towards the headend, then become combined together at the upstream receiver. The receiver has no way of knowing where this unwanted signal is getting in (ingress), and the noise impacts all cable modems on the same physical RF connection. This is particularly problematic because one ingress problem can impact an entire service group, consisting of hundreds of customers. Figure 6 shows the noise present, affecting a large service group. Finding the source of the ingress remains elusive and usually requires a maintenance technician. This is one of the primary motivators to keep our connectors tight.





**Figure 6 – Upstream Noise Impact on Entire Service Group**

Occasionally, some upstream noise problems are localizable to a small area, but these are relatively uncommon compared to the more typical upstream ingress. These types of problems can be caused by faulty return amplifiers which have a detectable distortion effect on the upstream signal for a smaller group of cable modems. With these “pocket issues,” it sometimes works to group devices that have a common RxMER or FEC-related problem. Unfortunately, this technique is minimally effective at localizing most typical upstream noise or ingress problems.

#### **4.1. The Troubleshooting Process**

In days prior to COVID-19, the troubleshooting process typically would begin with the technician making contact with the customer to obtain their perspective, in-person. This can transpire any number of ways and usually relies on customers using their own language to describe the service problem. For this to be effective, the technician needs to interpret the customer’s description and translate it to any number of troubleshooting and repair processes. For example, the customer may describe the problem as “the cable isn’t working,” which would require additional questions to further inform a) which cable service (video, voice, broadband) and b) which troubleshooting process to use. Next, the technician may ascertain that the customer was referring to internet service, then further determine if the connection is Wi-Fi or wired, and so on.

This is the input-side of the instruction loop, usually referred to as “features” in machine learning nomenclature. It’s easy to understand how translating the customer’s intent can result in misunderstood features. David Monnerat, Sr. Director of AI at Comcast, noted, “I’ve been on truck rolls in three states, all good people trying to do the right thing. They have good knowledge and skills but not a consistent

process, all troubleshooting differently. Three out of five technicians would have a different resolution to the same problem.” We’re faced with a wide variety of products, experience and technical understanding across our customers, agents and technicians. Fortunately, our tools and technology have been evolving to help bridge these divides and produce more consistent outcomes for our customers.

## **4.2. Repair Feedback**

After completing repairs, the technicians are usually required to provide some form of feedback about the work that was done to resolve the customer-reported problem. Typically, this is done when a technician closes a job (work order), and they are prompted to provide some mandatory selection from a list. This list can be vast and may allow for multiple selections, which will be discussed further. Most operators have used systems similar to this and experienced the same results regarding the unintentional bias in the resulting feedback. For instance, technicians might arbitrarily select the first code from a long list, regardless of the code.

### **4.2.1. Technician Performance Metrics**

Among of the most influential drivers of bias within the repair-feedback loop are the performance metrics used to measure technician productivity and effectiveness. These metrics are used for reports that assess how well technicians are doing, ultimately resulting in pay/career growth or the opposite, relative to their performance against the operator’s established goals.

Lessening repeat or re-work is a common goal when a service call requires multiple technician visits within a specified time period, such as a month. Given the many possible scenarios of service repair, there are loopholes which can insulate technicians against demerits from having to go back at a later time. Depending on the reporting algorithms, a technician may improperly report that the customer was not home, even though they may have performed some repair activity.

Another of the more common re-work loopholes is to complete the work order as avoidable (not required) by coding the repair with something like “no trouble found.” In this case, a technician may perform some repair activity, but depart unsure that it actually solved the problem. If they think there is a reasonable chance that the problem may be intermittent and unsolved, this provides some cover for a repeat visit in the future.

Although it doesn’t directly impact the repair coding of work orders, it’s not uncommon for technicians to leave their personal contact information with customers, so as to call them directly. This is another way for them to circumvent demerits associated with re-work. If the technician needs to go back and perform additional repairs or support, they do it off-record, avoiding negative performance reporting.

### **4.2.2. Customer Relationships**

The financial policies of cable operators can often influence technician repair feedback. It is not uncommon for operators to have a policy that establishes rules about when customers should pay for the service call versus it being free-of-charge, at the expense of the operator. These service call fees are often between \$50 and \$75, depending on the circumstances. Examples of chargeable service calls include scenarios such as a customer incorrectly re-connecting equipment when rearranging furniture, or if the power was simply turned at a power strip. In either of these examples, a minor oversight by the customer could prove embarrassing and is unfortunately all too common. Many technicians can empathize in these situations and find some menial form of repair that can be done to avoid an uncomfortable conversation about having to charge the customer for their oversight. It’s simple enough to code the work order as replacing a connector or splitter, versus having to break the bad news about a charge for the service call.

There are many areas where technicians have never charged a customer for a service call, providing anecdotal evidence that this is an all-too-common scenario.

#### ***4.2.3. The Expert Technician / Opportunistic Repairs***

There is a population of technicians that is especially enthusiastic about performing high-quality work. These technicians take pride in leaving every customer in better condition than before they arrived. They will scour the premises, starting at the tap, looking for loose or corroded connectors, then inspect the drop for leaks or damage, continuing to the ground block for distress or connector problems. These experts will invariably find plenty of opportunities for proactive repair, such as replacing old F-connectors, corroded ground blocks and imperfect cables. These are all great qualities that we hope all of our technicians would exhibit. However, they will often reflect every aspect of the repair in the work order. The multitude of proactive / opportunistic repairs may not have affected the original reason for the trouble call, although they certainly represent great hygiene for the network and insulation against a repeat trouble call. Also, by including multiple repair codes, this sometimes influences their performance statistics in a positive way. Naturally, this can result in an inflated number of repair codes that might not relate to the customer-reported problem.

#### ***4.2.4. Report Burden and Selection Bias***

Although it may not seem over burdensome to provide thorough feedback after a service call, the reporting process after a long, hard repair job can prove mentally exhausting – especially if the technician doesn’t believe the outcome of the report will have a meaningful effect on them or the business. In business terms, this represents a form of “decision fatigue,” resulting in the technician not putting sufficient energy towards making a proper selection to describe the job. This is exacerbated with growing lists, of hundreds of codes, which can take a long time to scroll, read and contemplate. Several controlled studies at Comcast have shown that the top code arbitrarily gets picked the most. Further attempts to randomize the top selection result in randomization of the repair disposition. The results were predictably consistent with the first presented code.

Some operators have attempted to reduce the selection burden by ordering the codes by their predominance. For example, replacing equipment, resetting devices and reprovisioning service are usually among the top selected codes. As a matter of improving the user experience, these codes might be ordered as the top three in the selection list, to reduce the searching and scrolling required by the technicians. Unfortunately, by placing the top selected code in first place, the previously discussed problem of arbitrary selection becomes compounded. This creates a selection loop bias which further strengthens the predominance of a small number of repair codes used to characterize a repair.

#### ***4.2.5. Intermittent Problems and Upstream Noise***

The nature of RF performance and troubleshooting can sometimes be intermittent. Upstream RF noise is notoriously spurious in nature. The source of the noise may be intermittent, such as turning on electrically noisy equipment, like the electric motors used in hair dryers or power tools. It’s also possible that the place the noise is getting in may be intermittent, such as a loose connection on a drop cable that might be blowing in the wind. As the wind blows, this can cause unpredictable shielding faults within the connector’s threads that intermittently allow the noise to enter the drop. If a customer calls at the time of the ingress and impaired service, it may be hours or days until a service technician may be able to visit the customer. It’s also possible that the ingress may be coming in at a completely different location on the node, unrelated to the customer who’s calling about the problem. In these cases, technicians will typically attempt some hygienic repairs, such as replacing connectors, without being able to confidently assert that

the underlying problem was fixed. In these examples, the technician may code the repair as no trouble found, replaced connectors, swapped equipment or any number of things.

#### **4.2.6. *Insufficient Tools or Training***

As our services become increasingly more complex, it can take time for our technicians to learn how to properly troubleshoot and repair them. For example, when Comcast introduced DOCSIS 3.1 service, there was a learning curve and tooling upgrades that were required to diagnose and repair orthogonal frequency division multiplexing (OFDM) signals. There was a time when legitimate RF problems could have been impacting a customer's experience, but the technician lacked proper tools or training to identify and repair the issue. There are other problems such as capacity, congestion, software updates and a myriad of others that may not be presented to the technician's troubleshooting process. The result of an undiagnosed problem typically results in the "Hail Mary" approach. Left with no other options, a technician will often fall back to replacing the customer equipment, otherwise known as a box swap. The vast majority of times, the equipment is not at fault, but it is possible that this activity does improve the service. At the very least, it offers the technician an opportunity to demonstrate that they are doing something perceived as helpful. In some cases, the act of swapping equipment usually includes re-provisioning, which can help. An example would be correcting an incorrectly-provisioned boot file which was unnoticed by the tools or technician. It could also be possible that a new device would provision to a different RF channel set which is less impaired than the previous implementation.

One of the most common forms of service repair is simply resetting or rebooting the equipment. While this does nothing to physically repair RF problems, it can sometimes be useful and has a very low cost, other than temporarily disrupting the service. In cases where the equipment software may be having problems, rebooting the device can temporarily re-initialize the software and restore proper function. However, this tends to be a temporary fix until the software bug is encountered again.

The latest versions of DOCSIS have proven to be exceptionally resilient. Our DOCSIS specifications have many coping mechanisms available to enable operation even in the most hostile RF environments. For example, a DOCSIS 3.1 cable modem may have 32 or more downstream channels (in addition to the OFDM signal) available for use. It is not uncommon to see frequency-specific problems, such as LTE ingress, which can impair a few channels, while the others might be operating perfectly. DOCSIS can bond different channel sets or disable problematic interfaces to allow error-free operation. This is another example of how modem resets can restore service without performing a repair to the physical environment.

## **5. Cloud Connected Technicians and Equipment**

Our field signal meters, and measurement systems have also been evolving, becoming more connected and integrated with the technician ecosystem. In our contemporary workforce, the signal analyzers are cloud-connected and augmented with all kinds of new information, such as technician identity, GPS coordinates, system design maps and telemetry measurements from customer equipment. This information-rich environment creates new opportunities for features and labels to provide to our ML and AI systems.

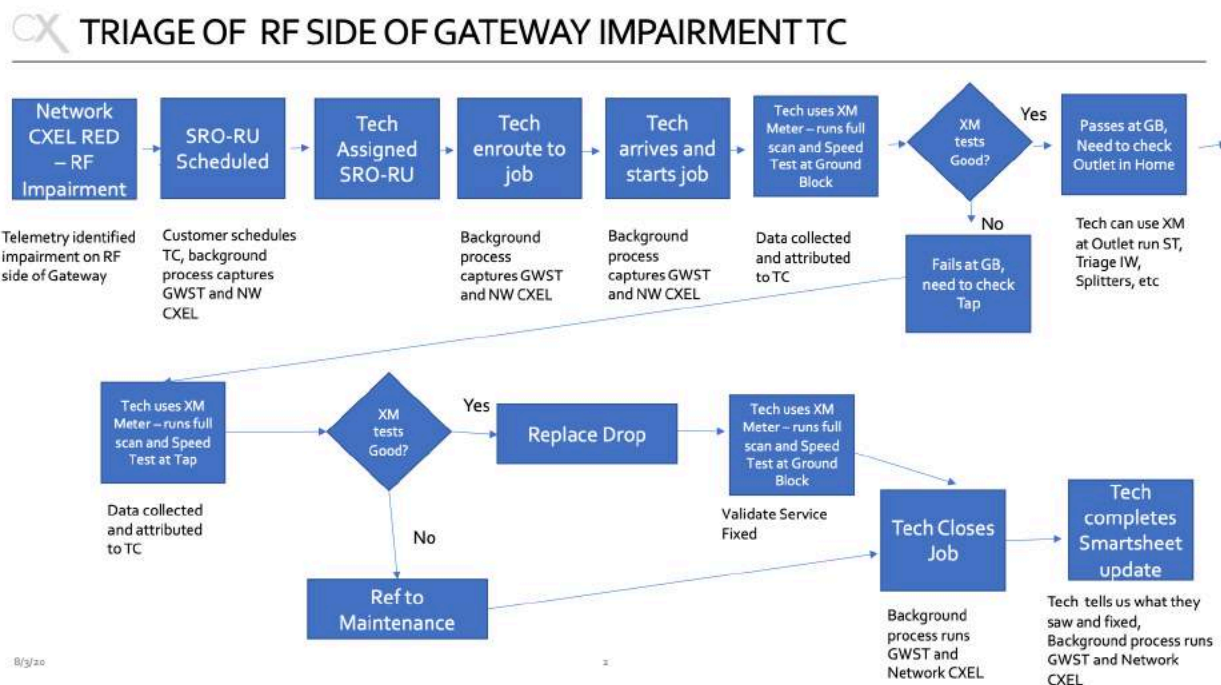
### **5.1. A New Troubleshooting Process**

As discussed previously regarding the troubleshooting process, it becomes obvious that this is an area of opportunity for improvement. With the onset of COVID-19, our customers, technicians, agents and

business partners now have an entirely new set of constraints and motivators shaping how they approach the troubleshooting process.

The first and most obvious constraint is the desire to maintain physical isolation between our technicians and customers when servicing their equipment. Brady Volpe, when asked about his experience on RF related repair calls, cites that “about 75% of the time it’s tap, drop or ground block. The other 25% of the time is in-home wiring ... these statistics exclude many of the common trouble call issues, such as customer education, and Wi-Fi problems, and varies by area.” Going by those statistics, and as a matter of efficiency, it makes sense to start the troubleshooting process outside. At Comcast, a new process was devised to help seize upon this statistical advantage and provide additional safety by allowing for physical separation. In addition to those two key benefits, it also facilitates a new opportunity for a consistent troubleshooting process.

As seen in the flow chart (Figure 7), technicians now start all RF troubleshooting by taking a signal measurement outside of the service location at the ground block. This is a critical demarcation, indicating where the cable service becomes physically attached and electrically bonded to a service location.



**Figure 7 – New Troubleshooting Process**

## 5.2. Process Compliance

Shortly after implementing cloud-based signal measurements, in early 2017, issues of process compliance became obvious. This was one of the first and most intriguing insights being provided by our new field measurement platform. A meager 22 percent of technicians were using their meters in a 24-hour period on any repair jobs. This quickly leads to all kinds of other questions about the troubleshooting process. How can someone possibly diagnose, repair and validate service without taking measurements? It stands to



reason that not all trouble calls require RF troubleshooting; however, these 22 percent of repair calls were service-related and resulted in RF specific repair codes.

At Comcast and a number of other operators, it has become common place to use the premises-located equipment to measure signal levels and performance. One of the leading causes for process non-compliance is that the technicians were relying on a premises health test (PHT) in lieu of taking measurements. This process evaluates the remote telemetry from the DOCSIS gateway and other equipment. While this is an acceptable way of validating service conditions, it's a procedural shortcut for the troubleshooting process. Lacking portability and segmentation of the network, this method results in a significant amount of "hunt and peck" rather than accurately diagnosing and repairing problems.

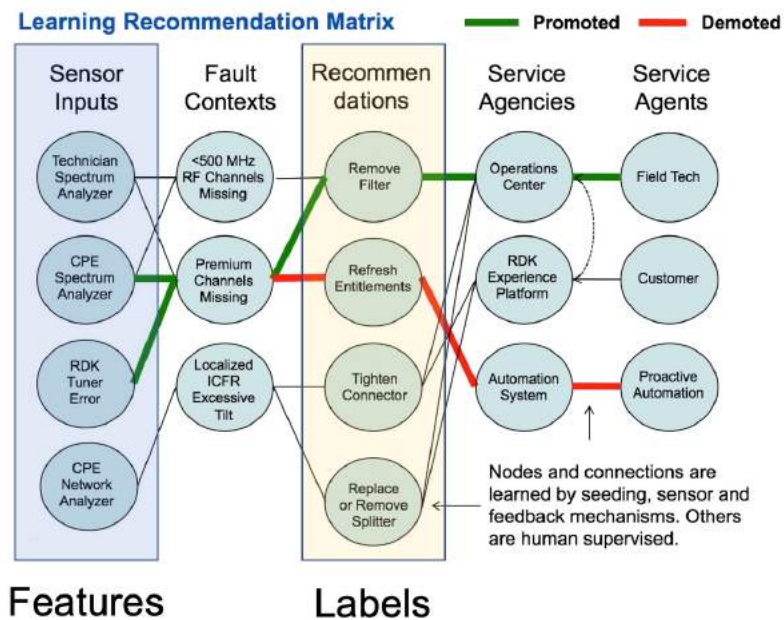
## **6. Machine Learning and Artificial Intelligence**

ML and AI have become ubiquitous in modern computing systems. Cable operators are investing heavily in these platforms and continue to find opportunities to improve how we operate our systems. This section reviews some of the techniques used, and results achieved when attempting to segment the RF network using ML and AI.

### **6.1. Features and Labels**

The review of the troubleshooting process was important to help convey an understanding of the features that we'll be using for our models. Features and labels are important constructs in ML and AI for both classification and regression problems. Features can be thought of as the inputs that will be available to the model. The features, in our case, will be any number of the measurement data which are available to our systems. Features might include in-channel frequency response (ICFR), full band capture (FBC) impairments, packet loss, speed test results, time of day and number of device resets.

Labels can be thought of as the output, or the desired prediction from the model. Examples of labels might be loose connectors, faulty drops or excessive splitters. These are the outcomes that we would expect our models to predict, given the proper features to inform its decision tree. Notice that the labels mentioned do not include customer behaviors, such as trouble calls or other service requests and interactions. One important aspect of this exercise is to decouple the objective service conditions from subjective customer experience. For instance, poor RxMER and packet loss may result in insufficient speed test performance. Some customers may call while others may not. That is discussed later.



**Figure 8 – Example of Cable-Oriented Features and Labels**

### 6.1.1. Machine Data vs Human Data

The previously discussed, fallibility of the human interpretation in our feedback system can become one of the most promising elements to improve our ML and AI model training potential. “You can think of machine learning using familiar terms to cable engineering. The goal of our models is to pick out the signal from the noise,” stated Jan Neuman, Executive Director of Machine Learning at Comcast. “Objective measurement that is noise-free and repeatable can be used to refine other measurements. By adding a less noisy signal, this increases the overall fidelity of the model. Thus, increasing the signal-to-noise ratio (SNR) results in more accurate predictions.” Put another way, improving prediction accuracy is the primary motivation for removing noise from our models.

Looking back on the discussion about the troubleshooting process, the first opportunity for noise begins at the input, or features. Although difficult to quantify empirically, the most fundamental features – such as “why did the customer call for help?”, or “which service is having a problem?” can be corrupt. If customer language is interpreted literally, “the cable isn’t working” could have many different meanings, causing noise in the feature. When this noisy feature is incorporated by the model and it starts making predictions, it is difficult to imagine that the results will be useful. By using machine data to refine the feature definition, the noise is reduced at the input of the model. An example of improving this feature with machine data would be a technician testing RF the ground block, outside of the location. If significant packet loss is measured, a noise-free feature now exists that is causal for a poor internet experience.

In addition to refining the features, machines offer an opportunity to improve the fidelity of the labels, or desired outcomes and predictions. While still not perfect, there are machine-provided data that can be reasonable proxies to some of the common labels. For example, the label of slow internet speed can be

approximated with automated speed testing. As expected, features of packet loss and labels of slow speed are clear machine signals that can be interpreted without added human noise.

### **6.1.2. Digital Interactions With our Customers**

For decades, our customers have been placing telephone calls to our call centers. As we drive towards more digital interactions with our customers, we’ve offered a unique and important opportunity to remove some degree of human interpretation (noise) from the models. Cable operators are embracing digital interfaces for our customers such as apps and online tools. By allowing customers to directly convey their intent, this offers an opportunity to bypass one of the most common causes of noisy labels, human-to-human communication on the telephone.

## **6.2. The Model**

Comcast Applied AI researchers developed a classification model to identify service calls where only outside of the premises work was required. The team developed and deployed this model to schedule trucks with a goal to minimize contact between technicians and customers during the COVID-19 isolation policy. After the initial model was deployed, they explored additional data, such as field meter measurements, to improve the model’s performance.

### **6.2.1. Outside of Home Applied AI Model**

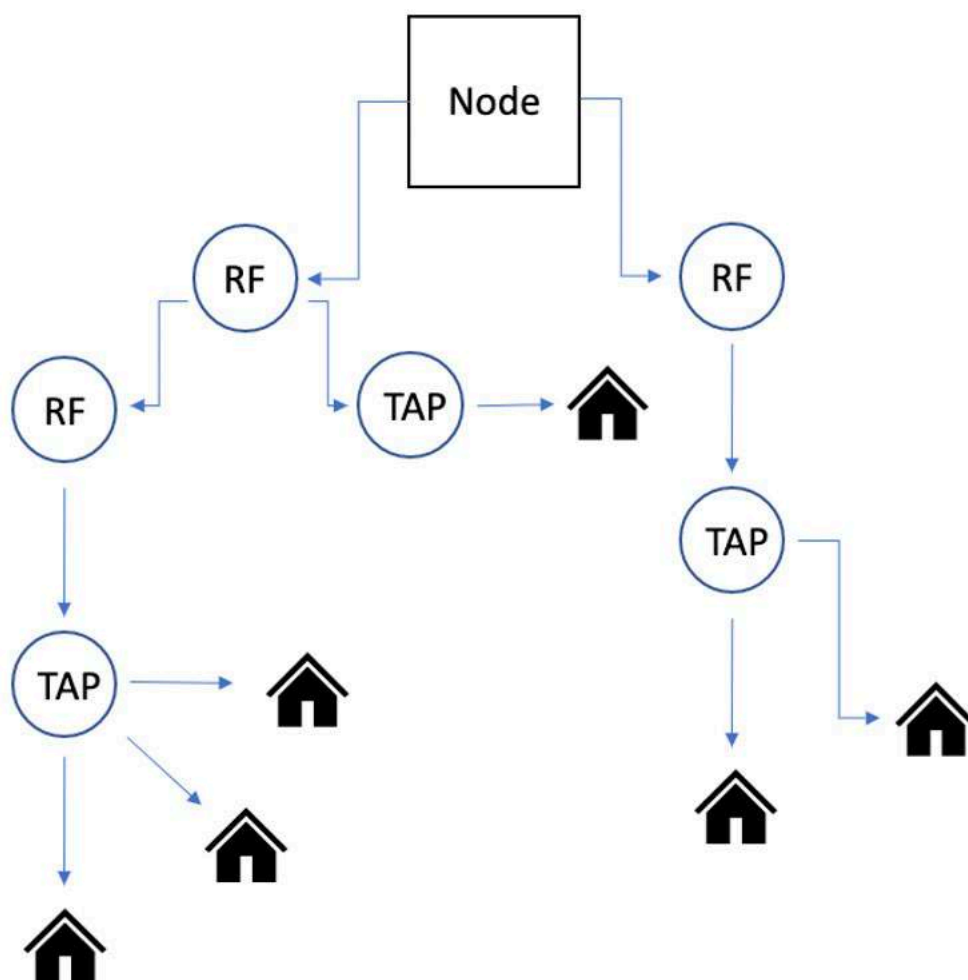
The “Applied AI outside of home” machine learning model used features from several network telemetry sources that collect and aggregate DOCSIS measurements. These sources poll and analyze the network, looking for outages and impairments. Other systems collect network data and do the fault segmentation analysis described earlier, which was also included. Table 1 describes the different data sources.

| <b>Source</b>                         | <b>Description</b>                                                                                                                                                                                                                                                                                          |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Account type                          | Account-level detail showing types of devices in the home (video gateways, wireless gateways, modems, etc.), days since account origination and since last device installation.                                                                                                                             |
| Prior truck history                   | Summary of technician-reported problems aggregated at the node level.                                                                                                                                                                                                                                       |
| Account network degradation algorithm | Features developed from an algorithmic tool that polls devices four to six times per day to report account-level degradation issues related to disturbances in the RF spectrum. Features are based on issue counts related to network, drop, in-home wiring, loose connections, and isolated home concerns. |
| Connectivity between modem and CMTS   | Features developed from a tool that polls devices three times per day to report raw measurements describing connectivity between the modem and the CMTS. Features include counts of impairment flags and means                                                                                              |

|                                   |                                                                                                                                                                                          |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                   | and standard deviations of measurements. Data reported include timeouts, system boot time, ripples, SNR, transmit and receive power, FEC errors, tap energy, and phase angle deviations. |
| Node network impairment algorithm | Features developed from an algorithmic tool that scans the network for RF impairments and reports them continuously. Impairments include plant Wi-Fi, suckout, wave, flux issues.        |

**Table 1 – Model Sources Summary**


Node topology was instrumental in creating features that detect network impairments. Node topology maps describe how equipment such as amps, taps, splitters, and cables are connected from the node down to customers' homes (Figure 9). Node topology maps were used to aggregate and average measurements, to provide a wider view of the network events surrounding a customer when the service call was being scheduled. One type of aggregation is a weighted average over a customer's node, where the weight is determined by the graph distance from the customer with the truck roll to other customers in the node. Another type of aggregation is at the parent level, which averages the telemetry over all customers who are immediate neighbors in the graph. We also computed averages for each piece of network equipment and compared it to averages for neighboring equipment, then found the piece of equipment that the customer depends on which has the greatest (and least) difference from its siblings. The idea is that if a single piece of network equipment is broken, we can observe this by comparing all the customers who depend on it (and thus have impaired service) with customers who do not depend on it but are otherwise similar because they share the same upstream network components.




**Figure 9 – Node Topology Schema**

Finally, the feature engineering pipeline aggregated impairments and telemetry data in a daily timeframe, predicting whether the truck roll would only require outside work.

The model supplied predictions to the agents, who advised customers on the type of truck to schedule. The team defined the target label using repair codes, provided by technicians after the service was completed. The outside fix codes included “refer to maintenance,” “construction,” or, for underground teams, “replace connector,” or “replace, repair, or run underground drop.” The best performing model was trained with an open source XGBoost classifier. The model was calibrated to achieve at least 5% recall and reported precision at 77% and lift at 2.48. Figure 10 illustrates the precision versus recall curve, normalized to a value of 1. Figure 11 shows the individual feature importance.



**Figure 10 – Model Performance, Precision and Recall**

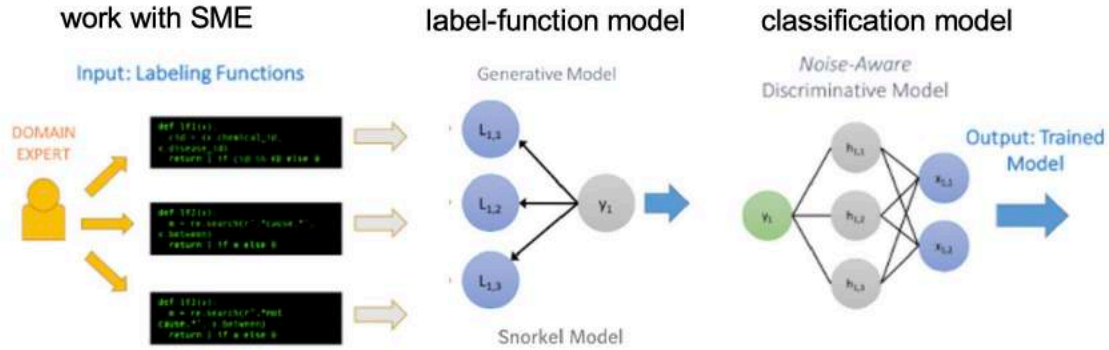


**Figure 11 — Model Feature Importance**

### 6.2.2. Using Weak Supervision to Improve the Model

After the model was deployed, the operations team launched a survey asking technicians to specify whether work was needed only outside of the house. These data were considered to be ground truth to verify the output predictions. Using these labels could potentially improve the model performance. However, ground-truth labels data were small and modern machine learning techniques require a large volume of data to train the model.

To address this problem, the team applied weak supervision techniques using the open-source Python package Snorkel (<https://www.snorkel.org/>). The goal was to access data not available at the time of prediction, use those data to develop a labeling model, train the labeling model on the ground truth data, and assign labels to the larger corpus of data that was lacking ground truth labels. Then, our data size would be sufficient to train a classification model on labels close to the ground truth and with sufficient data size. In other words, we attempted to use the results of signal meter readings after a small number of repairs, to predict what will happen on all trouble calls before the technician gets there. Figure 12 illustrates the weak supervision approach.



**Figure 12 – Weak Supervision Workflow**

The team gathered a corpus of data to use specifically for label development. Many of those sources were not available at the time of prediction. Those data included technician comments, repair codes, and signal meter data, among others.

Technicians took signal measurements on-site at the time of the truck visit. Of the total number of repairs, approximately 50% of them had field meter measurement data available. Engineers pre-processed the data for the team and developed labels by aggregating measurements such as per-channel RxMER, receive power and ingress. They assigned the following labels to each measurement: “ground block pass,” “ground block fail,” “tap pass,” “tap fail,” “refer to maintenance,” “tap,” and “tap fail ground block pass.” Accuracy of the functions derived from signal meter data ranged from 55% to 62%, surpassing the prevalence of outside labels in the survey data (at 45%). Table 2 enumerates the function performance.

| function                         | case    | coverage | overlaps | conflicts | accuracy |
|----------------------------------|---------|----------|----------|-----------|----------|
| lf_survey_outsidefittings        | outside | 5%       | 5%       | 4%        | 100%     |
| lf_problem_tap                   | outside | 11%      | 11%      | 7%        | 100%     |
| lf_survey_abletorepair           | outside | 33%      | 33%      | 24%       | 100%     |
| lf_problem_groundblock           | outside | 5%       | 5%       | 4%        | 100%     |
| lf_problem_drop                  | outside | 18%      | 18%      | 13%       | 88%      |
| lf_problem_refer_to_maintenance  | outside | 7%       | 7%       | 5%        | 82%      |
| lf_problem_refer_to_underground  | outside | 4%       | 4%       | 3%        | 80%      |
| lf_xm_overallresult_gbfail       | outside | 6%       | 6%       | 3%        | 62%      |
| lf_xm_overallresult_drop         | outside | 2%       | 2%       | 1%        | 61%      |
| lf_xm_outsidelabel               | outside | 16%      | 16%      | 9%        | 59%      |
| lf_xm_overallresult_gbpas        | inside  | 16%      | 16%      | 16%       | 59%      |
| lf_xm_overallresult_rtm          | outside | 4%       | 4%       | 2%        | 57%      |
| lf_xm_overallresult_tappas       | inside  | 12%      | 12%      | 12%       | 56%      |
| lf_xm_insidelabel                | inside  | 41%      | 41%      | 41%       | 55%      |
| lf_xm_overallresult_tapfailgbpas | outside | 3%       | 3%       | 2%        | 55%      |
| lf_survey_issuefixed             | outside | 54%      | 53%      | 40%       | 55%      |

**Table 2 – Labeling Functions with Highest Empirical Accuracy**

An important aspect of this exercise was the collaboration with data and field engineers to write functions and rules about the ground truth label. For example, if meter test interpretation was “ground block pass” or “tap pass,” the label model assigned the “inside of home” label. If test interpretation resulted in “drop,” for example, functions offered the “outside of home” label. Each function conveyed a proposed rule for labeling “inside of home” or “outside of home” case.

Many of the functions used technician-entered repair codes and free-form text comments. A set of these codes detecting tap and ground block issues had very high accuracy. Other repair codes varied in accuracy but made great contributions to the labeling model when they were combined with functions derived from signal measurements.



42 total functions were developed based on the combined labeling data corpus. Naturally, resulting functions varied in accuracy and coverage; some had labeling conflicts. We pruned the functions to include only those with higher accuracy. Then, we applied a majority voter strategy to assign predicted ground-truth labels and resolve function conflicts. In this strategy, cases where functions offered contradicting predictions were labeled neither inside nor outside in the training data.

Finally, a classification model was trained using Snorkel-labeled data. Both, benchmark and Snorkel-labeled models were trained on equal feature sets using the XGBoost classifier. We retrained both models on a limited set of data to benchmark the gains that can be achieved with the weak supervision method. Table 3 shows the precision at 5%+ recall target. The model trained with the snorkel-labeled data exceeded the performance of our benchmark model, with precision increase from 62% to 68%. By incorporating signal meter readings collected on-site, we were able to improve the accuracy of our labels and our model.

| Model                          | Precision | AUC   | Prevalence | Lift | Data size |
|--------------------------------|-----------|-------|------------|------|-----------|
| Ground truth label (benchmark) | 62.15%    | 0.545 | 45.8%      | 1.35 | 34,860    |
| Snorkel-labeled                | 68.32%    | .555  | 45.8%      | 1.49 | 370,154   |

**Table 3 – Model Performance Comparison**

## 7. Conclusion

Although all manner of misdirection (noise) exists within the data, there are still valuable insights to be gained. Our research has demonstrated that weak supervision techniques, access to subject matter experts, and a corpus of data are useful in developing labeling functions can help us improve the model performance. Specifically, access to diagnostic tools such field meter measurements can help us to improve training data labels.

Machine learning is a logical next-step toward identification and isolation of problems in the RF plant. Adding new data – in this case, our signal meter data – enhances RF domain expertise, improves operational performance and eliminates repeat service calls. By converging the ML/AI predictions with a real-time recommendation and feedback loop, there are also operational improvements to be realized (Figure 13). In other words, by using the results learned from highly effective technicians, the entire workforce can improve, resulting in more efficient technicians and improved customer experience.

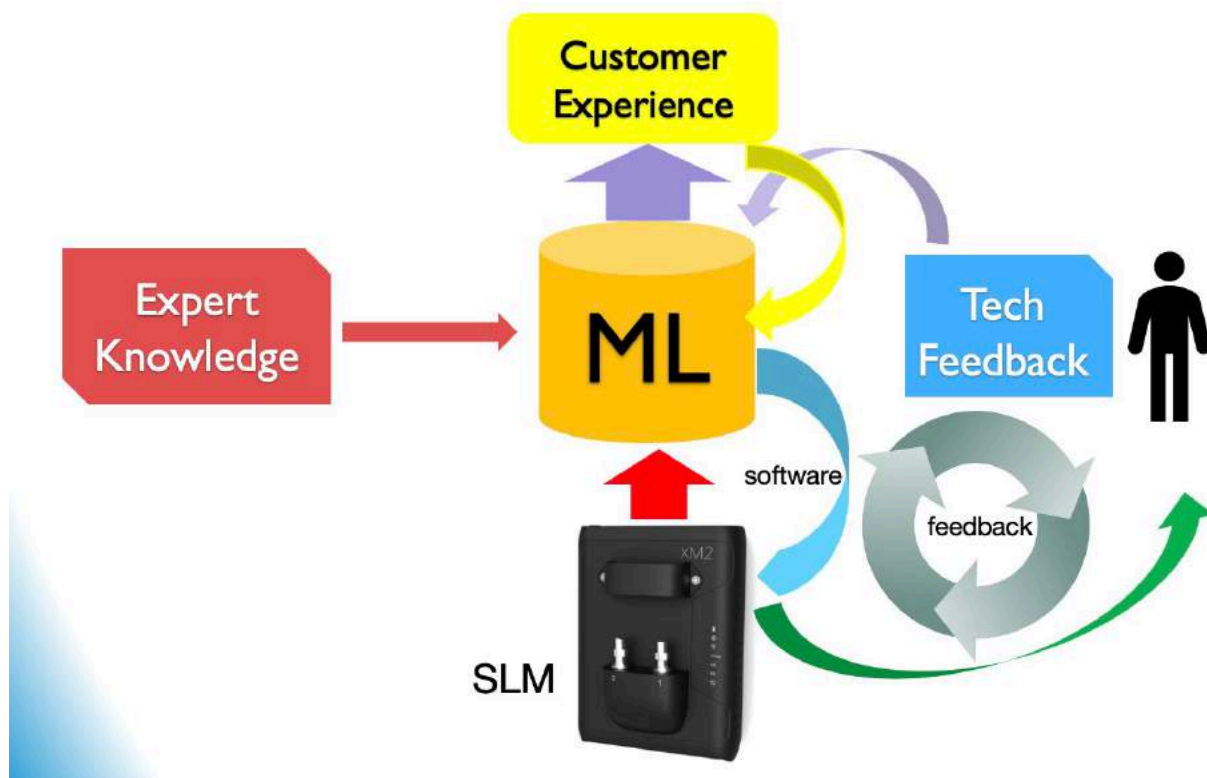


Figure 13 – ML/AI Enhanced Process and Feedback Loop

## Abbreviations

|        |                                                  |
|--------|--------------------------------------------------|
| AI     | artificial intelligence                          |
| CMTS   | cable modem termination system                   |
| DOCSIS | Data-Over-Cable Service Interface Specifications |
| FBC    | full band capture                                |
| FEC    | forward error correction                         |
| GPS    | global positioning system                        |
| ICFR   | in-channel frequency response                    |
| LTE    | long term evolution                              |
| ML     | machine learning                                 |
| OFDM   | orthogonal frequency division multiplexing       |
| PHT    | premises health test                             |
| PNM    | proactive network maintenance                    |
| RF     | radio frequency                                  |
| RxMER  | receive modulation error ratio                   |
| SID    | spectral impairment detection                    |
| SLM    | signal level meter                               |
| SNR    | signal-to-noise ratio                            |

## Bibliography & References

- [1] CableLabs DOCSIS® Best Practices and Guidelines - Proactive Network Maintenance Using Pre-equalization. 2012
- [2] SCTE Cable-Tec Expo 2016 - A Comprehensive Case Study of Proactive Network Maintenance; Wolcott, et al
- [3] Snorkel - A Weak Supervision System, May 31, 2019; Shreya Ghelani

# **Exposing The Invisible Enemy**

## **How Network Location Intelligence and Analytics Saves Lives**

A Technical Paper prepared for SCTE•ISBE by

**Sameh Yamany**  
CTO

VIAVI Solutions  
Sameh.yamany@viavisolutions.com

**Paul Gowans**

Director Solutions Marketing  
VIAVI Solutions  
Paul.gowans@viavisolutions.com

# Table of Contents

| <b>Title</b>                                              | <b>Page Number</b> |
|-----------------------------------------------------------|--------------------|
| 1. Introduction.....                                      | 3                  |
| 2. The MobileTelecommunications Revolution.....           | 3                  |
| 3. Mobile Network Advantages.....                         | 4                  |
| 4. The Role of Location Intelligence .....                | 4                  |
| 5. The Privacy Debate.....                                | 5                  |
| 6. Better Informed Decision Making.....                   | 5                  |
| 7. Real-Time Pandemic Hotspots Tracking And Analysis..... | 6                  |
| 8. Ensure Enacted Policies Make a Difference .....        | 6                  |
| 9. The Day After.....                                     | 6                  |
| 10. Conclusion: A Future with No Pandemics.....           | 7                  |
| Bibliography & References .....                           | 7                  |

## List of Figures

| <b>Title</b>                                                               | <b>Page Number</b> |
|----------------------------------------------------------------------------|--------------------|
| Figure 1 - Example City Scape with Different verticals and locations ..... | 3                  |
| Figure 2 - Comparison of App Based and Network Based Location Data.....    | 4                  |
| Figure 3 - Deriving Insights from Mobility Data .....                      | 5                  |

## 1. Introduction

A flu pandemic swept the world from 1918 to 1920, infecting approximately a quarter of the world's population and killing over 50 million. A century later, the world is again facing the reality of a viral pandemic. But things are different this time around: we have accumulated 100 years of health and technological advances that can fortify the fight against this pandemic. The invisible, hard-to-track, rapid transmission and replication of today's virus a critical danger, and these characteristics have caused the world economic and social wheels to grind to a halt. Absent a broader toolkit against this virus, the default defense strategy is isolation with everyone going on lockdown within their homes. Fortunately, now there are mature technologies in the world's arsenal that are being used to fight, in a smart way, the viral transmission and containment dilemma. These technologies are saving lives while enabling the gradual return to some social and economic normalcy. We will continue to embrace technology as the pandemic evolves, and to help with any future outbreaks.



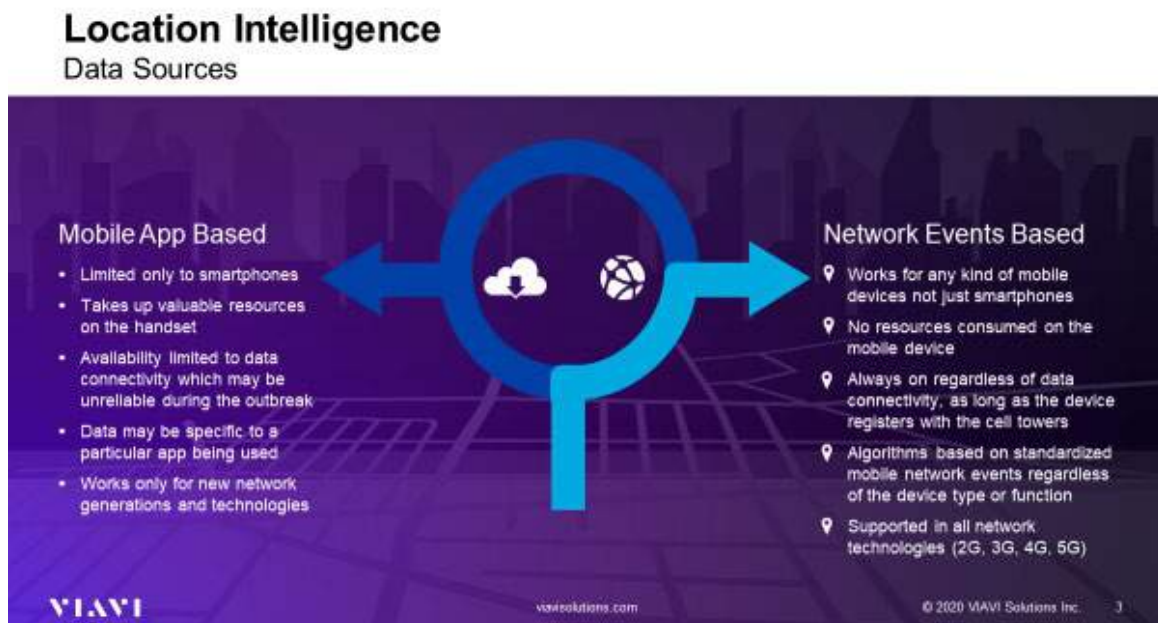
**Figure 1 - Example City Scope with Different verticals and locations**

## 2. The MobileTelecommunications Revolution

A 2019 Pew Research Center Fact Sheet <sup>1</sup> showed that almost 96% of the US population owned a cell phone of some kind, up from 35% in 2011. Almost Two Thirds <sup>2</sup> of the world's population today have mobile phones, and this percentage is above 90% if you only consider the top industrial nations. Compare this to the few million switchboard-based telephones during the 1918 flu pandemic. In the early 1970's the first handheld cell phone was invented. Fast forward fifty years, and more people have mobile phones than toilets <sup>3</sup>. Astonishingly, if we create a 2 minute time-lapse video showing the fast and explosive penetration of mobile technologies in the last thirty years, and compared it with a 2 minute-time-lapse video showing the new virus infections rate throughout the world in the first three months of 2020, the similarities between the two videos would be eye opening.

### 3. Mobile Network Advantages

Unlike the hard-to-track location of new virus transmissions, all mobile phones rely on mature network technology that manages, tracks and operates these devices in real-time to provide several services and advantages. Many countries are using their mobile network to help track, surveil, and stem the spread of the novel virus pandemic. The mobile network offers a level of data and location-based analytics using anonymized, highly granular, and accurate data and location intelligence insights that are being used by local and regional decision makers and health professionals to save lives and control the spread of the virus. There are also significant advantages to using network based location data versus application-based location data. Indeed, there has been significant issues in countries using app-based data rather than network based data to help track and trace<sup>4</sup>. Testing of Apps has shown issues with the technology at scale, despite Apple and Google working together on a common platform. Using network based (subscriber generated) data is well known and deployed in many operators world-wide covering all network technologies and always-on regardless of data connectivity issues.



**Figure 2 - Comparison of App Based and Network Based Location Data**

### 4. The Role of Location Intelligence

At VIAVI, we are working closely with communications carriers and governments around the world developing and deploying solutions that leverage Location Intelligence (LI) and Machine Learning (ML) technologies derived from mobile network events. These solutions, traditionally used for engineering, optimization and troubleshooting, provide powerful, accurate and reliable geolocation of subscriber's activities across the network. Correlating these activities with other core network usage insights turns the entire mobile network into a data analytics engine. Applying Machine Learning (ML) and data mining algorithms on such an engine provides valuable and actionable insights into network traffic levels and patterns of population movement and usage behaviors. In addition, detecting international roaming and inter-states movements within the network helps to accurately identify where the majority of subscribers are and where they've been. These insights can be tracked back in time for weeks, and in some cases, several months, all depending on the amount of historic data that has been stored. With 5G



communications, Location Intelligence will apply to much denser networks, both vertically and horizontally, adding another 3D dimension to deliver insight into network activities.



**Figure 3 - Deriving Insights from Mobility Data**

## 5. The Privacy Debate

The issue of privacy is often raised when looking at the role of Location Intelligence because many countries have very strict regulatory policies on the use of this type of data. Operators have used individual subscriber data for troubleshooting and diagnostics at the consent of the subscriber. Here, Location Intelligence is predominantly anonymized and aggregated to look at patterns, trends, and changes rather than individual analysis. The use and availability of this data is generally dependent upon the specific country and its regulations.

The process of anonymization is critical. It needs to be performed in a manner that both protects the privacy of the individual end user while allowing the mobile operator or the relevant governmental agencies to perform the relevant analysis. So far, it has been most effective to use a 24-hour static ID as the anonymized ID for each user in order to be able to track the mobility and determine if this represents reduced mobility, for example, in the society at large. If the same ID is used for a shorter period of time, critical mobility may not be observed. And if the ID is used for a longer period of time, it becomes possible to analyze who is behind the anonymized ID.

## 6. Better Informed Decision Making

Location Intelligence analytics enables decision makers to track travelers to and from affected countries and determine when and where quarantine rules can be enforced. The algorithms use the Location Intelligence aggregate data over time to identify hotspots that at-risk subscribers visited or where they congregated to detect potential contamination areas and implement proactive containment policies. The solution also monitors, in real-time, events (lawful or unlawful) for potential outbreaks and enforcement issues.

It also assists in outbreak investigations to identify contacts and apply appropriate measures to prevent further spread. This is particularly important as borders start to open up where countries go through different waves at different times. Being able to capture analysis on hotspots (e.g. where and when the



population has been moving across borders) becomes more important. Indeed, as life gets back to something that looks like normal, location data and analytics can help to manage outbreaks as it can continuously monitor people movements.

## **7. Real-Time Pandemic Hotspots Tracking And Analysis**

While the virus is invisibly attacking communities through travelers and locals carrying it from highly infected areas, its effect can still be uncovered, tracked, and quickly contained using Location Intelligence analytics. The Location Intelligence technology tracks inbound roamers from affected countries or high-risk areas and monitors their locations and interactions in real-time. It also alerts outbound roamers and provides the latest information on the outbreak and government policies to keep the roamer away from harm and potential hotspots. We continue to work with several governmental departments for disease control along with operators that requested help to monitor and manage the spread of the pandemic in their respective countries. In cooperation with in-region telecom operators, we are providing them with Location Intelligence and Machine Learning systems analyzing location data for inbound roamers and at-risk or positively-diagnosed cases of the virus. The system automatically identifies at-risk inbound populations, creates logical geofencing of quarantined patients with automatic alerting, and provides movement history and contact history for positive cases allowing notification and action to be taken to minimize infection risk. It also monitors critical infrastructure including hospitals, airports, and emergency services to assure they are getting the high connectivity, capacity, and performance from the network.

## **8. Ensure Enacted Policies Make a Difference**

As many of us enter the next phase in defeating the pandemic and gradually restoring life to normal, the ability to monitor and ensure the policies enacted—such as shelter in place, work from home, and curfews—is making a difference in the war against the virus. Here again, Location Intelligence and Machine Learning solutions can be of great assistance. Our solution helps governments and telecom carriers across the world with implementation of their policies by providing them with insights and actionable information that tracks potential contamination areas and reinforces verifications, alerts and health support for these policies. Again, this is done by monitoring patterns and movements of people, analyzing transport usage (location can determine speed and so can make predictions on whether someone is for example walking or in a motorized vehicle).

As more and more governments order school closures, and as people are encouraged to work from home, an immediate load and shift in network traffic was noticed for collaboration applications on the operators' networks. Our solution helps to detect this shift in real-time and recommends load-balancing and optimization action to mitigate these issues and assure efficient use of network and services resources.

Working from home policies may continue to be the norm for many people as they embrace a more blended model of work place and home. This shift may be permanent with operators needing to continue to analyze usage by location enabling networks to adapt and change to the evolving environment and delivering an optimized network at all times.

## **9. The Day After**

The use of Location Intelligence and Machine Learning does not stop when a pandemic ends. The question remains of how to deal with the aftermath—the unforeseen challenges that will arise from the social tensions, the economic impact, the shortage of essentials and food, unemployment, and so forth. The evolution of solutions like Location Intelligence, Network Assurance, and Machine Learning to

Artificial Intelligence (AI) helps governments and authorities to be prepared to tackle these issues with real time actionable information and decision making. Remote proactive monitoring and automation to actuate changes in the network or the services will be needed to offset the expected personnel shortages and the wide-spread nature of events

## 10. Conclusion: A Future with No Pandemics

Will we be able to completely defeat and eradicate viral pandemics in the next 100 years? The answer may be yes, if we consider the exponential advances in Artificial Intelligence, Machine Learning, data analytics, and the Internet of Things (IoT) coupled with the promises of 5G and next-generation networks. That conclusion is not a big stretch if we extrapolate the advancements achieved in the 100 years since the last deadly viral pandemic. The use of smart technologies is saving lives today, and the upward curve of innovations is not slowing. The network of the future will have the capability to detect a new viral pandemic at its onset, track all the infections in real-time, and stop it before spreading. Fully automated testing and in vivo monitoring are leaving the science fiction realm to become practical applications. Ultimately, we are on the side of technological innovations and human ingenuity to win the war against pandemics.

## Bibliography & References

1. *Pew Research Center Mobile Fact Sheet, June 2019* <https://www.pewresearch.org/internet/fact-sheet/mobile/>
2. *Statista: Number of mobile phone users worldwide from 2015 to 2020* <https://www.statista.com/statistics/274774/forecast-of-mobile-phone-users-worldwide/>
3. *Forbes* <https://www.forbes.com/sites/timworstall/2013/03/23/more-people-have-mobile-phones-than-toilets/#7c3aa9216569>
4. *BBC News, June 2020* <https://www.bbc.co.uk/news/technology-53114251>

# **MLaaS Applications in Digital Video – Supplanting Disliked Content**

A Technical Paper prepared for SCTE•ISBE by

**Srilal Weerasinghe PhD**  
Principal Engineer  
Charter Communications  
8560 Upland Drive, Englewood, CO 80112-7138  
720-699-5079  
Srilal.weera@charter.com

# Table of Contents

| <b>Title</b>                                                             | <b>Page Number</b> |
|--------------------------------------------------------------------------|--------------------|
| Introduction .....                                                       | 3                  |
| Content .....                                                            | 3                  |
| 1. Machine Learning Applications in Digital Video.....                   | 3                  |
| 2. Recommender Systems.....                                              | 4                  |
| 2.1. Matrix Factorization .....                                          | 4                  |
| 2.2. Singular Value Decomposition (SVD).....                             | 5                  |
| 3. Supplanting Unappealing TV Programs.....                              | 7                  |
| 3.1. Alternate Content Usage.....                                        | 7                  |
| 3.2. Content Replacement – Process Steps.....                            | 8                  |
| 4. Replacing Disliked Content – ML Automation .....                      | 9                  |
| 4.1. Challenges in Applying Recommender Systems to Disliked content..... | 9                  |
| 4.2. Implicit Identification of Disliked content.....                    | 9                  |
| 4.3. Data Rescaling .....                                                | 10                 |
| 4.4. Enhancements to Recommender Systems.....                            | 10                 |
| 5. Machine Learning in Digital Video – Additional Examples.....          | 11                 |
| 5.1. Enhanced Rating System for Movies .....                             | 11                 |
| 5.2. Video Content Analysis - VOD Storage and Ingest.....                | 11                 |
| 5.3. Personalized Ads - Combine Demographic and Viewing data .....       | 12                 |
| 5.4. Codec Quantization Parameter (QP) settings.....                     | 12                 |
| 6. Benefits to Service Providers and Programmers.....                    | 12                 |
| Conclusion .....                                                         | 12                 |
| Abbreviations.....                                                       | 13                 |
| Acknowledgements .....                                                   | 13                 |
| Bibliography & References .....                                          | 13                 |

## List of Figures

| <b>Title</b>                                                              | <b>Page Number</b> |
|---------------------------------------------------------------------------|--------------------|
| Figure 1 – Cosine Similarity between multi-dimensional vectors.....       | 4                  |
| Figure 2 – User-Item Rating Matrix .....                                  | 5                  |
| Figure 3 – Matrix Factorization.....                                      | 5                  |
| Figure 4 – SVD Decomposition.....                                         | 6                  |
| Figure 5 – SVD Decomposition in Elements Form.....                        | 6                  |
| Figure 6 – Process Flow for Alt Content Usage in Program Replacement..... | 8                  |
| Figure 7 – Disliked Content Usage to Enhance Consumer Profile .....       | 11                 |

## List of Tables

| <b>Title</b>                                                 | <b>Page Number</b> |
|--------------------------------------------------------------|--------------------|
| Table 1 – Alt Content Examples (names slightly changed)..... | 7                  |
| Table 2 – Disliked Content Rescaling.....                    | 10                 |

# Introduction

Machine Learning as a service (MLaaS) is a burgeoning field in the digital TV space. Its goal is to create AI/ML based revenue generating products. In this study, a novel use case is presented along with machine learning based enhancements. TV viewers routinely encounter shows that they dislike, but they are unable to avoid seeing them. While the consumer opinions are highly subjective, the end-result is the same: flipping the channel, which leads to advertising revenue loss for the programmer. Although retaining viewership of the channel is highly desired, technical challenges have precluded a satisfactory solution thus far.

The selected use case is of interest because unappealing content and recommendations contrast each other (dissuade vs. persuade). This distinction also manifests in the solution structure. For example, Recommender Systems (RS) are based on user ratings of liked content. In contrast, ‘disliked content’ may be so aversive to a viewer thus it is not even rated. Not having user ratings is a barrier for applying the RS model, which uses similarity measures in the latent space to determine affinity. Hence, in this study a different metric based on implicit data is used for feature vector creation. The goal is to illustrate the challenges and opportunities in developing MLaaS products for carrier-grade video.

Presented is a distributed solution\* applicable to vMVPD service. Enhancements to IP content delivery pipeline and Machine learning based automation are key for replacing disliked content. Additional scopes for MLaaS applications are also discussed.

\*patent filing (16/167,766)

## Content

### 1. Machine Learning Applications in Digital Video

Machine Learning applications in the video delivery pipeline are ubiquitous. From content ingest to transmission to delivery, opportunities abound for applying algorithmic solutions. These would typically include ingest quality control, network and storage optimizing and a host of data analytics applications upon content delivery. For example, the manual scanning of thousands of TV ads at ingest can be automated with a classification engine [1]. Additional applications in operational and product improvement are discussed later.

Such applications in cable-tech however are internal to the enterprise. The premise of this paper is to make the case that the technology is ripe for the next stage of ML revolution. Known as MLaaS or Machine Learning as a Service, it is modeled similar to other ‘as a service’ paradigms such as SaaS and PaaS. Familiar examples are web-based emails and cloud DVR services. MLaaS goes a step further and facilitates machine learning based consumer technology applications. As an enabling technology it can create machine learning based revenue generating services for the cable operators.

In the ensuing sections we delve deeper into the problem. First, we review the functioning of general recommender systems. Then, to illustrate MLaaS, a novel application is presented along with machine learning enhancements. The selected use case is advertising revenue loss due to channel surfing. The recommender systems algorithms are applied with certain caveats. Finally, other MLaaS applications are briefly discussed.

## 2. Recommender Systems

Consumers are overwhelmed by the broad array of choices in product availability. Personalized recommendations are thus vital for a more satisfying user experience. Recommendation and Search are both similar in that they attempt to filter relevant information from a deluge of data. In the case of search, the user knows what (s)he is looking for. Recommender Systems (RS) on the other hand, attempt to determine user interests via automated search analytics. The ‘items’ in this context can be books, movies etc. Machine learning algorithms enable learning from the data and building predictive models.

Recommender Systems generally fall into two broad categories: Content-Based filtering (CB) and Collaborative-Filtering (CF). The CB approach is based on the attributes of each item and the user’s affinity (rating) for similar items. Content may consist of movie genre, actors, theme etc. Based on the users previous ratings, movies that have similar content are recommended.

Unlike CB, the CF recommendations are based on multiple users’ ratings. The assumption is that similar users share the same interest. Sometimes called ‘neighborhood methods’, they are categorized as user-based or item-based. The user-based CF make recommendations based on the ratings given by other users with similar profiles. In item-based CF, the neighborhood items are those with similar user ratings. The similarity is determined by first forming vectors (rows of ratings matrix) per each user, and computing similarity measures such as Euclidean distance or Cosine similarity between vectors in a multi-dimensional space. Mathematically, users in the same ‘neighborhood’ would have Cosine of angle close to 1. Dissimilarity is denoted by 90° separation.

$$\text{similarity} = \cos(\theta) = \frac{\mathbf{u} \cdot \mathbf{v}}{\|\mathbf{u}\| \|\mathbf{v}\|} = \frac{\sum_{i=1}^n u_i v_i}{\sqrt{\sum_{i=1}^n u_i^2} \sqrt{\sum_{i=1}^n v_i^2}}$$

**Figure 1 – Cosine Similarity between multi-dimensional vectors**

(Technically there is a difference between Cosine Similarity and Cosine Distance due to Schwartz inequality. In this paper however, we follow the conventional definition and use inner angle as the similarity measure between the vectors.

### 2.1. Matrix Factorization

Among the algorithmic approaches to solve RS, matrix factorization is a well-known technique. The basic premise of matrix factorization is that there are *latent factors* that determine the user ratings of items. These latent factors are not measured directly, but their impact is reflected in the user ratings (observed variable). In the case of movies, those could range from obvious features such as action, romance and comedy to more complex psychological emotions. Such nuances may not have simple labels to describe them. The beauty of the matrix factorization method is that it can handle such complexity with poise via a multi-dimensional vector model.

The starting point is the user-item ‘rating matrix’ (utility matrix), which in practice may have millions of entries. Note that some cells are empty (user ratings not available).

|        | Movie 1 | Movie 2 | Movie 3 | ..... | Movie n |
|--------|---------|---------|---------|-------|---------|
| User 1 | 5       | N/A     | 2       |       | 3       |
| User 2 | 1       | 3       | N/A     |       | N/A     |
| User 3 | N/A     | 4       | 1       |       | 4       |
| .....  |         |         |         |       |         |
| .....  |         |         |         |       |         |
| User m | N/A     | 4       | 1       |       | N/A     |

**Figure 2 – User-Item Rating Matrix**

Matrix factorization decomposes the user ratings matrix based on latent factors that contribute to user preferences and item attributes. It then predicts the unknown ratings through the scalar/inner/dot product of the latent features of users and items.

For example, if Bob is an action movie aficionado, his profile may consist of:


Bob = 60% Action + 30% Comedy + 10% Romance+ 0% Historical

Similarly, the profile for movie could be:

Titanic = 30% Action + 0% Comedy + 60% Romance + 10% Historical

The dot product of these vectors would determine Bob’s affinity for the movie Titanic. The larger the dot product between a user vector and an item vector, the item is better suited for user’s taste and can be recommended.


The rating matrix of Figure 2 has many empty cells (sparse), as only a fraction of movies are watched by users, let alone rated. The goal of recommender system is to fill these gaps. In Figure 3, the sparse matrix is expressed in terms of two latent factor matrices (one for user and the other for item attributes). The missing values are approximated by the dot product of the two latent factor matrices via optimization.



**Figure 3 – Matrix Factorization**

## 2.2. Singular Value Decomposition (SVD)

A prominent technique for matrix factorization is singular value decomposition. The SVD algorithm received wide recognition for its critical role in the \$1 million Netflix prize competition [2]. The singular value decomposition (SVD) factorizes a given matrix ‘X’ into constituent arrays of feature vectors.



**Figure 4 – SVD Decomposition**

$$A = U\Sigma V^T$$

- $A$  is the given data matrix ( $m \times n$ ) of rank  $r$
- $U$  is an orthonormal  $m \times r$  matrix
- $\Sigma$  is a diagonal  $r \times r$  matrix
- $V$  is an orthonormal  $n \times r$  matrix

The diagonal matrix ' $\Sigma$ ' contains the singular values in descending order. The higher values represent the dominant features or latent factors of the given matrix. We consider the reduced form with small Sigma values set to zero as shown in the Figure 4.

Collaborative filtering (CF) based recommender systems perform poorly when dimensions in data increases (sometimes called 'curse of dimensionality'). Reducing the extra dimensions while keeping the salient features is the function of a dimensionality reduction algorithm. SVD filters the dominant features from the data (Sigma matrix) and identify hidden correlations in the singular vectors  $U$  and  $V^T$ .

SVD helps find a lower rank matrix approximation to the original matrix. Mathematically this means picking only the top singular values from the sigma matrix and discarding others. Each singular value defines the 'strength' of the concept. e.g. a high value may indicate the genre of the movie cluster such as action or comedy.

$$A \approx U_{m \times r} \Sigma_{r \times r} V_{r \times n}^T$$

$$\begin{pmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots \\ a_{m1} & \dots & \dots & a_{mn} \end{pmatrix} \approx \begin{bmatrix} u_{11} & u_{12} & \dots & u_{1r} \\ u_{21} & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots \\ u_{m1} & \dots & \dots & u_{mr} \end{bmatrix} \begin{bmatrix} \sigma_{11} & 0 & \dots & \dots \\ 0 & \sigma_{22} & \dots & \dots \\ \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \sigma_{rr} \end{bmatrix} \begin{bmatrix} v_{11} & v_{12} & \dots & v_{1n} \\ v_{21} & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots \\ v_{r1} & \dots & \dots & v_{rn} \end{bmatrix}$$

**Figure 5 – SVD Decomposition in Elements Form**

An alternate way to perform matrix factorization is to calculate eigenvalues of the covariance matrix ( $A^T A$ ). Both approaches are consistent in that the square roots of eigenvalues are equal to singular values. The Eigen vectors of the covariance matrix indicate 'principal components' of the matrix. Eigen method (sometimes called spectral decomposition), is computationally more involved, hence in this paper we explore SVD.



As mentioned, the main challenge with the matrix factorization methods is that the rating matrix is mostly sparse. This problem is addressed by iteratively calculating each cell value and updating them (using Gradient Descent and Alternate Least Square techniques). The error is minimized by applying Frobenius norm to data matrices. These steps are well-known in the literature and as such not discussed further.

We revisit the recommender systems in Section 4, in reference to disliked content.

### 3. Supplanting Unappealing TV Programs

In this section we change gears and review the technical aspects of the selected use case. Then, the machine learning enhancements are discussed for creating an MLaaS product.

Almost all TV viewers have their favorite shows, as well as the ones they dislike. The latter (unappealing content) is the focus of this study. When confronted with disliked content the normal user behavior is to flip the channel, which leads to advertising revenue loss for the programmer. IP based digital TV streaming has made it possible to supply alternate content on per user basis.

While IP streaming offers such capability, the requirements cannot be met by changing just one end of the content stream. The contractual obligations between Programmer and Content Distributer must remain intact as well. A distributed solution is presented, requiring enhancements on multiple components: Alternate content is supplied by the programmer, assembled and stored at the Content Distributor. The disliked content is replaced automatically.

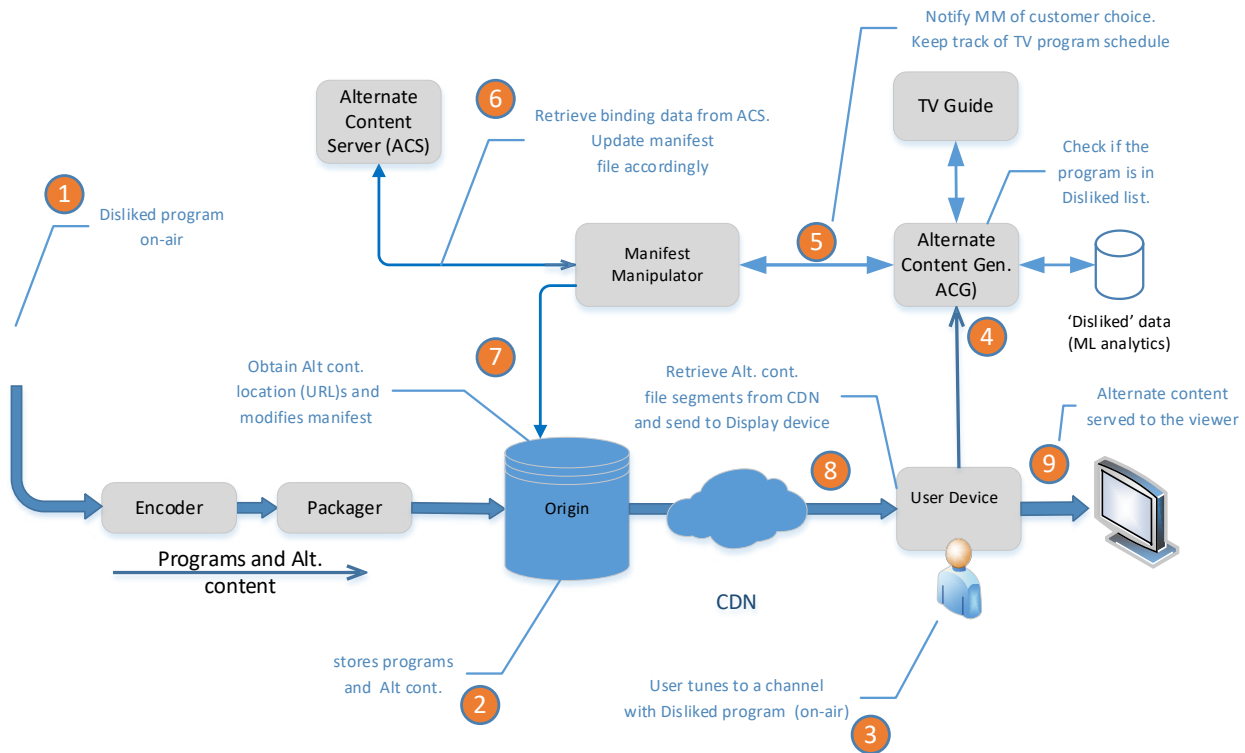
#### 3.1. Alternate Content Usage

A customer identifies a TV program/show/series that is unappealing to her, and would prefer to have it replaced with alternate content. Her choices can be supplied via a web portal/API or a clickable icon added to the TV Guide. The Content Provider establishes Alternate content channels for selected regular channels. The Alternate content could be programs that aired earlier in the day and stored in the CDN. Once the setup is complete, whenever the consumer tunes to an unappealing program content, it will trigger content replacement automatically and seamlessly. In this context, *'tuned to the channel'* could also mean accessing a web page, such as in a social-media based network.

**Table 1 – Alt Content Examples (names slightly changed)**

| Supplanting a Program by Alt Content (Traveler Channel) |                |                  |                  |
|---------------------------------------------------------|----------------|------------------|------------------|
| Original Program                                        | Alt content-1  | Alt content-2    | Alt content-3    |
| Ghoul Adventures                                        | Bizarre Feasts | Museum Mysteries | Nay Reservations |

### 3.2. Content Replacement – Process Steps



**Figure 6 – Process Flow for Alt Content Usage in Program Replacement**

1. Assume a previously identified ‘disliked program’ is on the air (e.g. Table 2 - *Ghoul Adventures* of Traveler Channel).

2. Program content and alternate content are stored on the CDN headend. The Alt. content could be previously aired shows or specifically targeted content.

(In a typical streaming scenario - The broadcasted TV/media content stream from Programmer is received by the content distributor. The Encoder/Transcoder may perform any format changes. The Packager/Segmenter will splice the content into many chunks. The video/audio segments (along with an Index file (Manifest) for segment identification), are placed for storage on a CDN Origin server).

3. Assume the previously identified ‘disliked program’ is on the air. The end-user tunes to the channel.

4. Alternate Content Generator (ACG) interfaces with TV guide to obtain program start end times. It queries the ‘Disliked content database’ for any matches.

5. If a match is found that information is passed back to the Manifest Manipulator. (Note – These messages are in addition to the regular manifest requests issued by a Client device during IP streaming)

6. Noting that there is a ‘Dislikes’ request, the Manifest Manipulator communicates with the Alternate Content Service (ACS) database. ACS module **binds** Customer Device ID to Program Alt. cont. selection and returns to the Manifest Manipulator.

Customer Device Identifier  $\leftrightarrow$  Disliked content  $\leftrightarrow$  Alt. content

7. Manifest Manipulator modifies the index/manifest file, accordingly, replacing regular content with alternate content. The updated manifest is sent to the customer device (Client).
8. The Customer Device now retrieves corresponding alternate content from the CDN and send to the display device.
9. The original content is supplanted with alternate content and supplied to the display device for user consumption.

## 4. Replacing Disliked Content – ML Automation

### 4.1. Challenges in Applying Recommender Systems to Disliked content

In the movie recommendations example described earlier, the viewer inputs were explicitly supplied (e.g. one through five stars). In the case of disliked content however, there is no explicit metric. The content is so aversive, the viewer simply changes the channel. Not having user ratings is a barrier for applying the RS model, which uses similarity measures in the latent space to determine affinity. Hence, in this study a different metric is used for feature vector creation.

Another constraint is the sheer number of channels in a TV subscription. Most users have never even tuned to the hundreds of channels offered. It would be incorrect to categorize those as disliked content. Note that the present analysis applies to channels containing a mix of liked-content and disliked-content. The goal is to keep the viewer in the same channel. If there are no liked programs at all (i.e. customer never tunes to the channel), then the issue of ad revenue loss is moot.

### 4.2. Implicit Identification of Disliked content

The method adopted was to collect viewership data over time and apply data analytics to identify disliked content indirectly. One peculiarity with the viewership data is the vast difference in the time scale. For ‘disliked content’, the channel surfing times are about one to two seconds. But for ‘liked content’ (regular viewing), the durations could vary from several minutes per channel to hours. This disparity was addressed via a change of scale adjustment.

The wide range of the time scale is also a barrier for applying machine learning algorithms. This is illustrated below with reference to two algorithmic scenarios.

#### a) Cosine Similarity

Cosine similarity is easy to visualize in two dimensions, but its application in the present context is for a multi-dimensional space. Due to the large disparity in the time scales (few seconds for disliked content vs. thousands of seconds for liked content), the user/item vectors in the latent space will not be an accurate depiction. The impact of disliked-content will be hard to quantify. On the other hand, simply inverting the scale (e.g. assigning 5 stars for disliked and 1 for liked), will not capture the nuances.

#### b) Singular Value Decomposition (SVD)

Note that in the Singular values matrix, the dominant values are at the top left and decreases down the diagonal.

$$\text{Sigma } 1 \geq \text{Sigma } 2 \geq \text{Sigma } 3 \dots$$

The quantities of interest (disliked content with 1-2 seconds), are at the lower end of the time-scale and their impact is washed out by much larger terms in the matrix (thousands of seconds for liked content). As such, the dominant Sigma values do not reflect the impact of disliked-content, and skew the results. It is therefore necessary to rescale the data so that the disliked content reflect the dominant terms in the utility matrix.

### 4.3. Data Rescaling

The metric adopted for this task is the log reciprocal of channel surf time. Since the channel change time could vary from 1 second to several hours (thousands of seconds), we use the following formula:

$$\text{Disliked Content Measure} = \text{Log}_{10} (10^4/T) = 4 - \text{Log}_{10} T$$

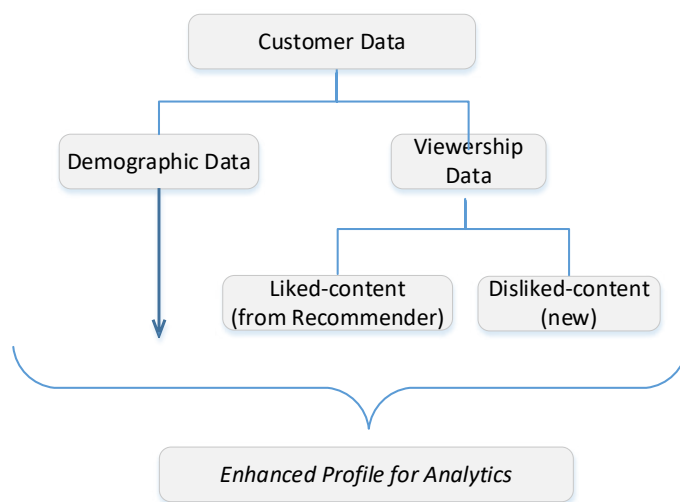
**Table 2 – Disliked Content Rescaling**

|               | ← Disliked Content |        | Liked Content → |         |
|---------------|--------------------|--------|-----------------|---------|
| Duration      | 1 sec              | 10 sec | 15 min          | 2 hours |
| Log10 (104/T) | 4                  | 3      | 1.045           | 0.143   |

The logarithmic scale converts the wide range of channel view time values to a more manageable compact scale. The original span of 1 second to over 2 hours, is now rescaled to the more compact range from zero to four (Table 2). The Log scale captures 1–5 star ratings implicitly. It is also in line with behavior based algorithms due to the continuity of scale. Conversely, simply assigning 5 stars to shorter times would not be granular. Also, the user response times for channel surfing are subjective.

### 4.4. Enhancements to Recommender Systems

The result of applying RS is the creation of a disliked content matrix. This complements the more common liked-content matrix created by RS. The disliked-content analysis supplies additional data about user preferences to enhance the recommender system. This would be useful since a major issue with RS is the sparsity of the input data matrix.



**Figure 7 – Disliked Content Usage to Enhance Consumer Profile**

## 5. Machine Learning in Digital Video – Additional Examples

Additional instances of potential usage are listed below for reference; not all are revenue related.

### 5.1. Enhanced Rating System for Movies

Current TV/Movie ratings are confusing (TV-Y7-FV, PG-13, TV-14...), as the restricted content definition is subjective. Parents would appreciate if they could know beforehand the type and placement of restricted content.

- a) Scan videos in the repository and tag restricted content per pre-defined criteria, with time stamps.
- b) Develop an API for one-click access for the 'Enhanced Ratings' from the program guide.

### 5.2. Video Content Analysis - VOD Storage and Ingest

- a) Scan videos and generate descriptive metadata. Identify sentiments, underlying topics as well as any anomalies in the media content. Create a searchable catalog of videos based on tagged data.
- b) Thematic advertising – Given an Ad-campaign theme (e.g. eco-tourism), find matching videos from the collection. Find effectiveness of ads by different demographics/audiences.
- c) Celebrities - Find videos of a given actor, including duration/time stamps, from a collection.
- d) Closed captions – Translate speech to text for assets that currently do not have captions.
- e) Skipping content on VOD – User is presented with the option to auto-skip parts of a video based on pre-defined content identifiers. ML can tag content based on *heuristics*.

### 5.3. Personalized Ads - Combine Demographic and Viewing data

[Scenario – Demographic data may suggest that 3 people live in a household, but it cannot answer questions such as, “Of the 3 inhabitants, which one was watching TV at 4 PM)?” Data analytics based on viewing patterns can be used to enhance targeted advertising]

- a) Develop individual user profiles for each household based on multiple data sources – customer demographics, online viewing history and navigation data.
- b) Make recommendations for personalized ads based on time-of-day viewership characteristics.
- c) Predict ad completion rates based on past viewing behavior and make ad recommendations.

### 5.4. Codec Quantization Parameter (QP) settings

In a video codec the DCT coefficients are quantized per Quantization Parameter (QP) settings. Coarse QP values mean high compression and lower quality. QPs are also proportional to the Lagrangian multipliers of Rate optimization.

The Neural Network based auto-encoders on the other hand, map the data to a lower dimensional latent space and then reconstruct it during decoding. It can predict pixel values fairly accurately and provide optimum QP values. The encode-decode pair is trained as a single unit in unsupervised learning. The adaptively tuned QP values will yield improved PSNR.

## 6. Benefits to Service Providers and Programmers

Machine learning enables identifying and auto-replacing of disliked content. MLaaS based novel offering (“Don’t Like, Don’t Watch!”) would be a new revenue opportunity for programmers and content distributors. The disliked content could even be faces and voices based on pre-defined signatures.

Another improvement is in targeted advertising. Ad campaigns are generally based on consumer demographic data obtained from data brokers. The aggregated data can be refined by comparing with viewership patterns (e.g. time-of-day), garnered from liked and disliked content (Figure 7).

Recommender systems based on explicit data can be improved by combining with disliked-content analysis.

## Conclusion

Cable-tech is ripe for disruption and transformation with MLaaS based consumer technologies. Service providers will reap the benefits of new revenue generating opportunities. To illustrate the point, a novel use case was presented along with MLaaS enhancements. To circumvent the challenges with general recommender system model, a new metric was proposed based on *implicit* user data. The auto-discovery and replacement of disliked content would prevent revenue loss for programmers due to channel surfing. Additional benefits include enhancements to recommender systems and targeted ad-campaigns.

## Abbreviations

|       |                                                                         |
|-------|-------------------------------------------------------------------------|
| AI/ML | Artificial Intelligence/Machine Learning                                |
| MLaaS | Machine learning as a service                                           |
| vMVPD | Virtual Multi-Channel Video Programming Distributor (Internet based TV) |
| SVD   | Singular Value Decomposition                                            |
| MF    | Matrix Factorization                                                    |
| CB    | Content Based Filtering                                                 |
| CF    | Collaborative Filtering                                                 |

## Acknowledgements

The author would like to thank Prof. Athula Gunawardena for helpful discussions.

## Bibliography & References

1. Srilal Weerasinghe, *Machine Learning Applications in Cable TV Advertising*, SCTE-Cable-Tech-Expo-2019
2. Yehuda Koren, *Netflix Grand Prize*,  
[https://www.netflixprize.com/assets/GrandPrize2009\\_BPC\\_BellKor.pdf](https://www.netflixprize.com/assets/GrandPrize2009_BPC_BellKor.pdf)

# **Proposal of RF/IP Adaptive Video Distribution Scheme over Cable Television Access Networks**

A Technical Paper prepared for SCTE•ISBE by

**Yoshitaka Kidani**

Research Engineer

Japanese Cable Laboratories

KDX Kayabacho Bldg. 3F, 3-4-2 Nihonbashi Kayabacho, Chuo-Ku, Tokyo 103-0025, Japan

+81 3 5614 6100

y-kidani@jclabs.or.jp

**Hiroyuki Yamashita**

Research Engineer

Japanese Cable Laboratories

h-yamashita@jclabs.or.jp

**Shuichi Matsumoto**

President

Japanese Cable Laboratories

s-matsumoto@jclabs.or.jp



# Table of Contents

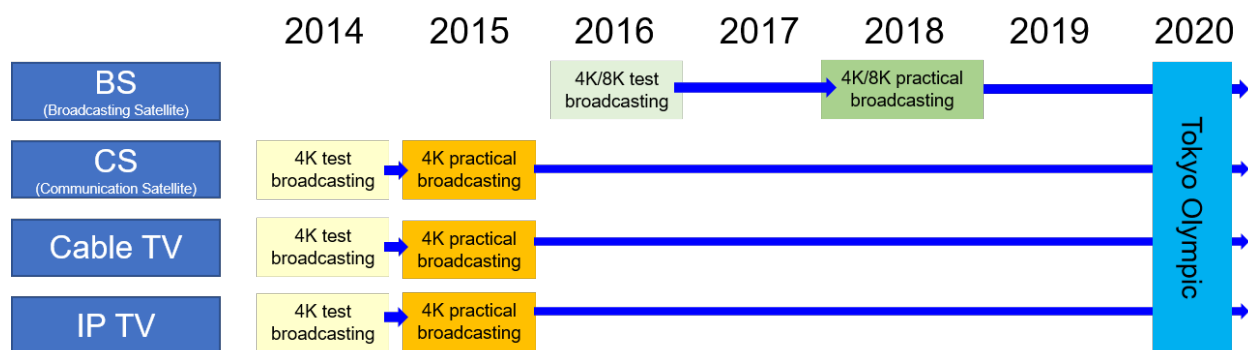
| Title                                                                         | Page Number |
|-------------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                          | 3           |
| 2. Problem statement .....                                                    | 3           |
| 3. Proposed method .....                                                      | 4           |
| 3.1. Basic Concept .....                                                      | 4           |
| 3.1. System architecture, Technology .....                                    | 5           |
| 3.1.1. Optimizer / Switching Indicator .....                                  | 6           |
| 3.1.2. Optimizer / Switching Algorithm .....                                  | 6           |
| 3.1.3. Data collector, etc. / Real-Time and Large-Scale Data Collection ..... | 7           |
| 3.1.1. Switcher /Seamless Switching .....                                     | 7           |
| 4. Simulation Experiment .....                                                | 7           |
| 4.1. Simulation System .....                                                  | 7           |
| 4.2. Simulation Condition .....                                               | 7           |
| 4.3. Simulation Result .....                                                  | 8           |
| 4.3.1. Validation from Time Transition .....                                  | 8           |
| 4.3.2. Validation for Various Condition.....                                  | 8           |
| 5. Prototype Development of Switching Demo System .....                       | 9           |
| 6. Conclusion.....                                                            | 10          |
| Abbreviations .....                                                           | 10          |
| References .....                                                              | 11          |

## List of Figures

| Title                                                                                                                                                             | Page Number |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| Figure 1 - 4K/8K roadmap based on an interim report from the study group for 4K/8K broadcasting at the Ministry of Internal Affairs and Communication (MIC) ..... | 3           |
| Figure 2 - Typical utilization of HFC network band .....                                                                                                          | 3           |
| Figure 3 - The portion of access network types of Japanese cable subscribers in 2017. ....                                                                        | 4           |
| Figure 4 - Conventional scheme of video distribution over HFC access network. ....                                                                                | 4           |
| Figure 5 - Priority rule of controlling video qualities and distribution schemes. ....                                                                            | 5           |
| Figure 6 - RF/IP adaptive distribution scheme.....                                                                                                                | 5           |
| Figure 7 - A system architecture to realize the proposed scheme.....                                                                                              | 6           |
| Figure 8 - Simulation system configuration.....                                                                                                                   | 7           |
| Figure 9 - The difference of audience satisfaction degree between with the proposed method (switching) and without the proposed method (no switching). ....       | 8           |
| Figure 10 – Comparison of ASD in different types of cable infrastructure.....                                                                                     | 9           |
| Figure 11 – Configuration of switching demo system .....                                                                                                          | 9           |
| Figure 12 - Overview of switching demo system .....                                                                                                               | 10          |

## 1. Introduction


UHD video streaming services have been launched worldwide. In Japan, 4K/8K broadcasting services has been promoting by “All Japan,” which are organized by the government, broadcasters, telecommunicators, and TV manufacturers since 2014. Figure 1 is “4K/8K roadmap” advocated by the Ministry of Internal Affairs and Communication (MIC) and satellite, cable TV, and IP TV are picked up as the access networks for 4K/8K broadcasting services in the roadmap. According to this roadmap, cable industry in Japan has been started 4K community channel service organized by themselves, named by “Cable 4K”, since 2015.



**Figure 1 - 4K/8K roadmap based on an interim report from the study group for 4K/8K broadcasting at the Ministry of Internal Affairs and Communication (MIC)**


## 2. Problem statement

Demand for 4K/8K video is increasing toward the Tokyo Olympic games, cable industry seeks to increase 4K video channels to respond the demand from customers. The final goal of Japanese cable industry is to convert over the 100 community channels with SD/HD resolution into 4K. The transmission capacity of HFC access networks is, however, limited by the existing TV services and Internet services, and the network band cannot be expanded due to the retransmission of satellite broadcasting ( 1~3GHz). Figure 2 shows the typical utilization of HFC network band in Japan. Cable operators needs to manage to deliver the over 100 broadcasting services and Internet services in the limited network band.




**Figure 2 - Typical utilization of HFC network band**

Furthermore, cable operators cannot easily replace HFC by FTTH because of the high installation cost of FTTH. Figure 3 shows the portion of access network types of Japanese cable subscribers in 2017 and indicate the fact that HFC is still dominant. Because of this, the cable operators cannot easily enhance the network infrastructure to increase the 4K channels.



**Figure 3 - The portion of access network types of Japanese cable subscribers in 2017.**

Figure 4 shows the conventional scheme of video distribution over HFC access network. In Japanese, the most of the cable operators deliver the all broadcasting services on RF-network band as shown in Figure 4. This is because that the RF distribution scheme is consider to be more reliable than IP distribution scheme, that is IP multicast scheme. The detail is as follows. In general, with RF-distribution scheme, the received video quality is stable since the video is transmitted at a fixed coding bitrate over quality-guaranteed network where are fixed transmission band is reserved for each video channels. The utilization efficiency of video of the network band resource is, however, low since they occupied by the videos regardless audience ratings, In contrast, with the IP distribution scheme, the received video quality is unstable since the video is transmitted at a variable or adaptive coding bitrate over best-effort network where the bandwidth is not reserved for reserved for each video channels. The utilization efficiency of network band resources, is however, high since they are used for the transmission of not only video but also Internet data. Whichever method is used, therefore, it is difficult to realize 4K conversion of all existing video channels aimed at by the cable industry, and a novel approach is required.



**Figure 4 - Conventional scheme of video distribution over HFC access network.**

### 3. Proposed method

#### 3.1. Basic Concept

As the solution of the problem described in section 2, RF/IP adaptive distribution scheme is proposed. The strategy of proposed method is to create the environment where almost all the subscribers can watch 4K videos. According the strategy, the following two key ideas are proposed. The first idea is sharing of

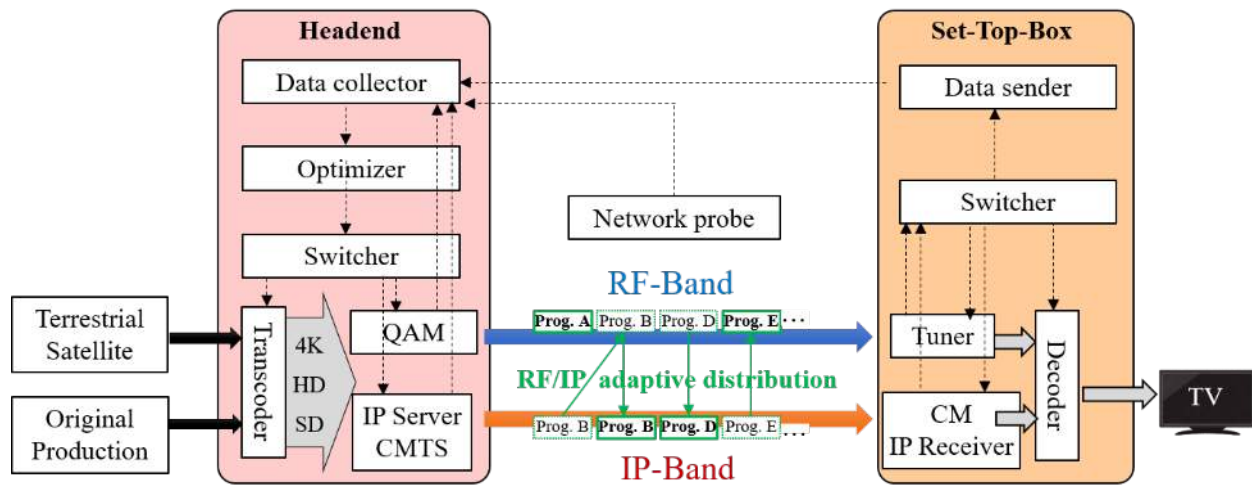
the RF/IP network band resources. The second idea is adaptive control of video qualities and distribution scheme according to the following three parameters; attribute of video content (e.g., emergency degree), available network bandwidth, and audience rating of content. Figure 5 shows the priority rule of video regarding the second idea. In this rule, the scheme quality, that is video resolution and bitrate we defined, is controlled according to audience ratings and emergency degrees. By preferentially distributing video channels with a high audience rating in 4K resolution and high bitrate, it is possible to realize an environment in which almost all subscribers can watch 4K channels. The reason why RF distribution scheme is higher priority than IP distribution scheme is that RF distribution scheme is more stable than that of IP described section 2.

**Figure 5 - Priority rule of controlling video qualities and distribution schemes.**

Figure 6 shows the proposed RF/IP adaptive distribution scheme. The basic principle is as follows: First, the optimizer collects the audience ratings, available network bandwidth, and attribute of video content. Second, the optimizer calculates the switching indicator and determines the necessity of switching the video distribution scheme for each video channel. Third, the optimizer signals a switching order to the switcher for each video channel which the optimizer determined that the switching is required. Forth, the switcher change the video distribution scheme according to the switching order from the optimizer.

### Figure 6 - RF/IP adaptive distribution scheme

In this section, a system architecture and key technologies of key components are described. Figure 7 indicate a system architecture to realize the proposed RF/IP adaptive distribution scheme. Data collector in addition to optimizer and switcher is also key component. The role and technology for each key component is described separately in the following section.



**Figure 7 - A system architecture to realize the proposed scheme.**

### 3.1.1. Optimizer / Switching Indicator

The switching indicator is designed as like QoE with the following formula. We call it audience satisfaction degree (ASD). The definition of this is the ratio of A and B as shown in the formula. A is a measured score calculated with sum of products based on the actual audience rating and the score corresponding to the bitrate of the actually used video distribution scheme. B is an ideal score calculated with sum of products based on the actual audience rating and maximum score which is corresponding to the bitrate of RF/4K video distribution scheme.

$$ASD(t) = \frac{A}{B} = \frac{\sum_{i=1}^n R_{t,i} \times S_i}{\sum_{i=1}^n R_{t,i} \times S_{max}} \times 100$$

It is noted that  $R$ ,  $S$ ,  $i$ ,  $n$ , and  $t$  indicate audience rating, score based on video bitrates, indicator of programme, total number of programmes, and distribution time, respectively.

### 3.1.2. Optimizer / Switching Algorithm

The switching algorithm is designed to keep always the highest score of ASD even when the audience ratings and available network bandwidth, and attribute of video content vary. It is realized by recursively checking and assignment based on the priority rules. The overall processing order is as follows:

1. Collect the latest ratings, bandwidth, attribute
2. Sort programmes by ratings
3. Assign programmes with high ratings to the RF-band at a high bitrate preferentially.
4. Lower the rank of the scheme and quality if bandwidth becomes insufficient
5. Recurse No.3 and No.4 processes until all programmes are completely assigned.
6. Calculate difference of new/old ASD based on the new/old assignments
7. Switch scheme and quality if the difference is greater than threshold

Some simulation experiments for validation of the switching algorithm are conducted and those are described in section X.

### 3.1.3. Data collector, etc. / Real-Time and Large-Scale Data Collection

Real-time and large-scale data collection in data collector, network probe, and CPE devices is essential to guarantee the accuracy of the switching scheme. To realized that, lightness and ease to implementation, will be required. From our preliminary study, MQTT (Message Queuing Telemetry Transport) is the best protocol in terms of the requirements.

### 3.1.1. Switcher /Seamless Switching

Seamless switching is essential to avoid subjective side-effect from the switching shock. These are the requirement:

1. How to transmit the switching timing from HE to STB
2. How to remove the switching shock (e.g., stop playback, skip frame, etc.)
3. How to solve the difference in video format (e.g., resolution, frame rate, codec)

Feasibility study of seamless RF/IP switching is conducted thorough the prototype development of switching demo system, which is described in section X.

## 4. Simulation Experiment

### 4.1. Simulation System

To validate the proposed switching algorithm, simulation system is developed with two web application servers. Utilized CPU and software are as follows:

- CPU: 1vCPU, Memory: 2GB, OS: CentOS
- Software: Apache Tomcat, mysql, openssl, Echarts

A simulation engine where switching runs is implemented from scratch. Figure 8 shows the simulation system configuration. Time-varying audience ratings and IP network bandwidths are input as scenario files. Other basic parameters are set with GUI before the simulation.

Time transition of ASD ( = switching indicator ) is plotted as a result of simulation execution.




Figure 8 - Simulation system configuration

### 4.2. Simulation Condition

Utilized basic parameters and time-vary parameters are shown as follows:

- Number of video channel: 50ch, 75ch, 100ch
- Network capacity: Utilize 8 patterns of infrastructure referred to Japanese operators
- Video coding bitrate:

- RF-4K: 25Mbps, RF-HD: 15Mbps, RF-SD:6Mbps, IP-SD: 4Mbps
- Number of STB: 10,000
- Emulation Time: 24 hours
- Switching Cycle: 5 min
- Switching Threshold: 0
- Time-varying parameters
- Audience rating: Utilize real data from Video Research Ltd.
- Internet traffic (affect the capacity of IP-bandwidth): Utilize statistical data of time transition of Internet data from MIC


### 4.3. Simulation Result

#### 4.3.1. Validation from Time Transition

Figure 9 shows the difference of audience satisfaction degree (ASD) between with the proposed method (w/ switching) and without the proposed method (w/o switching) where the number of video channels are 100 and network capacity is set as follows:

- Total of RF-network bandwidth: 900Mbps
- Total of IP-network bandwidth: 320Mbps

The red line and black line indicate the ASD with the proposed method and the ASD without the proposed method, respectively. From this result, it is confirmed that the proposed method has the effect to maintain high ASD value in all hours.



**Figure 9 - The difference of audience satisfaction degree between with the proposed method (switching) and without the proposed method (no switching).**

#### 4.3.2. Validation for Various Condition

The simulation of various condition regarding infrastructure type and the number of video channels are conducted. Figure 10 shows the comparison table of that. In total, 8 patterns of infrastructure types and 3 patterns of the number of video channels are tested as described in Figure 10. From this table, it is observed that ASD with switching is over double score of that without switching as the maximum effect of the proposed method under certain conditions (described as red square frame in Figure 10). This means that the proposed method has the effect to doubles the utilization efficiency of bandwidth or the number of 4K content compared to without this scheme.

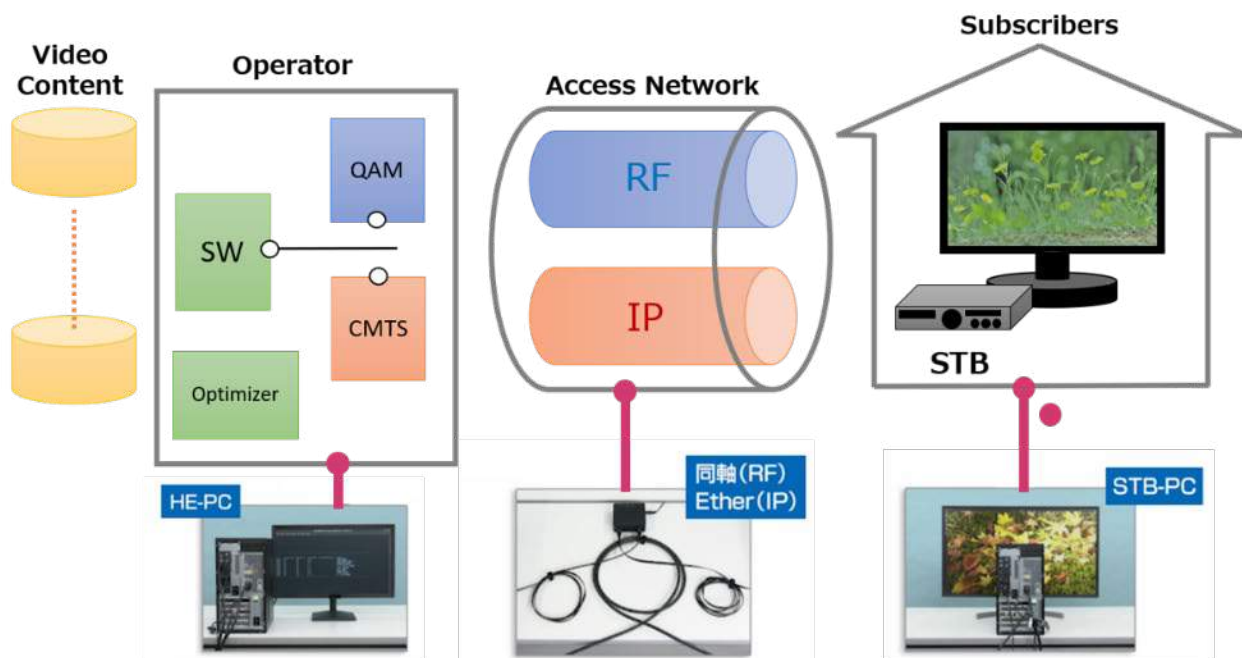


| Number of video channel (ch) →<br>↓ Infrastructure pattern (bps) |                                  | 50ch<br>= 1250M (if all channels are 4K)           | 75ch<br>= 1875M (if all channels are 4K)           | 100ch<br>= 2500M (if all channels are 4K)        |
|------------------------------------------------------------------|----------------------------------|----------------------------------------------------|----------------------------------------------------|--------------------------------------------------|
| 1                                                                | RF-band: 2070M<br>IP-band : 160M | N/A (Because all channel can be transmitted as 4K) | N/A (Because all channel can be transmitted as 4K) | w/o Switching: 83.7<br><b>w/ Switching: 99.0</b> |
| 2                                                                | RF-band: 1770M<br>IP-band : 160M | N/A (Because all channel can be transmitted as 4K) | w/o Switching: 95.4<br><b>w/ Switching: 99.9</b>   | w/o Switching: 72.3<br><b>w/ Switching: 97.6</b> |
| 3                                                                | RF-band: 1320M<br>IP-band : 160M | N/A (Because all channel can be transmitted as 4K) | w/o Switching: 73.7<br><b>w/ Switching: 97.7</b>   | w/o Switching: 52.7<br><b>w/ Switching: 94.3</b> |
| 4                                                                | RF-band: 1020M<br>IP-band : 160M | w/o Switching: 82.5<br><b>w/ Switching: 98.9</b>   | w/o Switching: 58.6<br><b>w/ Switching: 95.3</b>   | w/o Switching: 40.2<br><b>w/ Switching: 90.5</b> |
| 5                                                                | RF-band: 1950M<br>IP-band : 320M | N/A (Because all channel can be transmitted as 4K) | N/A (Because all channel can be transmitted as 4K) | w/o Switching: 80.2<br><b>w/ Switching: 98.6</b> |
| 6                                                                | RF-band: 1650M<br>IP-band : 320M | N/A (Because all channel can be transmitted as 4K) | w/o Switching: 90.6<br><b>w/ Switching: 99.5</b>   | w/o Switching: 69.7<br><b>w/ Switching: 96.9</b> |
| 7                                                                | RF-band: 1200M<br>IP-band : 320M | w/o Switching: 96.6<br><b>w/ Switching: 99.9</b>   | w/o Switching: 69.9<br><b>w/ Switching: 96.9</b>   | w/o Switching: 49.6<br><b>w/ Switching: 93.4</b> |
| 8                                                                | RF-band: 900M<br>IP-band : 320M  | w/o Switching: 78.5<br><b>w/ Switching: 98.2</b>   | w/o Switching: 54.4<br><b>w/ Switching: 94.7</b>   | w/o Switching: 36.5<br><b>w/ Switching: 90.3</b> |

**Figure 10 – Comparison of ASD in different types of cable infrastructure**


## 5. Prototype Development of Switching Demo System

To study the feasibility of seamless RF/IP switching, PC-based demo system is developed. Figure 11 and Figure 12 show the configuration and overview of switching demo system, respectively. We confirmed that the switching has been completed within 1 second from the switching order of the server PC (HE-PC) to switching video quality of client PC (STB-PC). It is noted that we used the 4K video test materials of ITE in the verification test.



**Figure 11 – Configuration of switching demo system**





**Figure 12 - Overview of switching demo system**

## 6. Conclusion

Highly-efficient video distribution scheme over the existing HFC network was proposed. Basic concepts are sharing the RF/IP resources and switching RF/IP distribution scheme according to audience ratings, network bandwidth, attribute of video content. From the simulation results and the prototype development of demo system, it was revealed that theoretical effectiveness and feasibility of the proposed scheme. Especially, experiment results show that proposed scheme doubles the utilization efficiency of bandwidth or the number of 4K content compared to without this scheme. The switching demo system show the evidence of feasibility for seamless RF/IP switching. A total feasibility study by implementing optimizer and data collector on the demo system, and international standardization are future works.

## Abbreviations

|        |                                                             |
|--------|-------------------------------------------------------------|
| bps    | bits per second                                             |
| CMTS   | cable modem termination system                              |
| CPE    | customer premises equipment                                 |
| CPU    | computer processing unit                                    |
| DOCSIS | Data Over Cable Service Interface Specification             |
| FEC    | forward error correction                                    |
| FTTH   | Fiber to the home                                           |
| GB     | giga byte                                                   |
| GUI    | graphical user interface                                    |
| HE     | headend                                                     |
| HFC    | hybrid fiber-coax                                           |
| HD     | high definition                                             |
| Hz     | hertz                                                       |
| ITE    | The Institute of Image Information and Television Engineers |
| PC     | personal computer                                           |
| IP     | internet protocol                                           |

|     |                                 |
|-----|---------------------------------|
| QAM | quadrature amplitude modulation |
| QoE | quality of experience           |
| RF  | radio frequency                 |
| STB | set-top-box                     |
| SD  | standard definition             |
| TV  | television                      |
| UHD | ultra-high definition           |

## References

ITE 4K Test Materials (<https://www.ite.or.jp/content/test-materials/>)

# **Unleashing Managed SD-WAN With Closed-Loop Automation**

## **A Standards Based Approach to Increased Operational Efficiency and Network Agility**

A Technical Paper prepared for SCTE•ISBE by

**Tom DiMicelli**

Senior Product Marketing Advisor

Ciena Corporation

7035 Ridge Road Hanover, Maryland 21076

410-694-5700

[tdimicel@ciena.com](mailto:tdimicel@ciena.com)

# Table of Contents

| <b>Title</b>                                                       | <b>Page Number</b> |
|--------------------------------------------------------------------|--------------------|
| Table of Contents .....                                            | 2                  |
| List of Figures.....                                               | 2                  |
| 1. Introduction.....                                               | 3                  |
| 2. SD-WAN Market Trends and Forecasts.....                         | 3                  |
| 3. Managed SD-WAN Service Evolution.....                           | 4                  |
| 4. Managed SD-WAN using Legacy Operations.....                     | 5                  |
| 5. The Business Impact of Legacy Operations .....                  | 6                  |
| 6. SD-WAN-related Standardization Efforts .....                    | 6                  |
| 6.1. MEF .....                                                     | 6                  |
| 6.2. TM Forum.....                                                 | 7                  |
| 7. Automating SD-WAN operations.....                               | 8                  |
| 8. Standards-based Managed SD-WAN Automation.....                  | 9                  |
| 9. The Business Impact of Standardized, Automated Operations ..... | 10                 |
| 10. Managed SD-WAN Automation in Action .....                      | 11                 |
| 11. Conclusion.....                                                | 13                 |
| Abbreviations .....                                                | 15                 |
| Bibliography & References.....                                     | 15                 |

# List of Figures

| <b>Title</b>                                                                          | <b>Page Number</b> |
|---------------------------------------------------------------------------------------|--------------------|
| Figure 1 – Enterprise SD-WAN adoption (source IDG) .....                              | 3                  |
| Figure 2 – SD-WAN management preferences, source IDG.....                             | 4                  |
| Figure 3 – Manual SD-WAN service and network configuration.....                       | 5                  |
| Figure 4 – MEF LSO Reference Architecture diagram .....                               | 7                  |
| Figure 5 – TM Forum Open Digital Architecture framework .....                         | 8                  |
| Figure 6 – SD-WAN Automation framework .....                                          | 10                 |
| Figure 7 – Automating Network as a Service Catalyst reference diagram .....           | 11                 |
| Figure 8 – Windstream Enterprise legacy operational workflow for SD-WAN Services..... | 13                 |
| Figure 9 - Windstream Enterprise's automated SD-WAN service activation workflow.....  | 13                 |

## 1. Introduction

Managed SD-WAN services underpin many strategic enterprise initiatives and provide cable operators with a high-growth market. As the market matures and cable operators expand their SD-WAN service portfolio, they must broaden their service portfolios and improve their operational processes to ensure they maximize this opportunity.

This paper summarizes SD-WAN trends and growth forecasts and highlights both market and operational challenges associated with managed SD-WAN implementations today; it also describes how progress that has been made in industry standards bodies can help cable operators implement an open, standards-based closed-loop automation framework that overcome these issues.

## 2. SD-WAN Market Trends and Forecasts

Enterprises are embracing SD-WAN as the foundation of their strategic SaaS, cloud, and digital transformation initiatives. The 2019 SD-WAN Market Trends Survey<sup>1</sup> showed that 54 percent of survey respondents had already deployed SD-WAN, and that enterprise adoption could reach 90 percent.





Figure 1 – Enterprise SD-WAN adoption (source IDG)

The results of this survey, and forecasts from other analysts, indicate that the SD-WAN market will likely grow at double digit rates for the foreseeable future. In fact, even though Vertical Systems Group lowered their 2020 U.S. managed SD-WAN revenue growth outlook due to pandemic-related disruptions, they also forecast high-growth rates through 2024,<sup>2</sup> as SD-WAN's ability to provide efficient connectivity to cloud-hosted applications are well-suited to post-pandemic business needs.

<sup>1</sup> IDG MarketPulse Research, 2019 SD-WAN Market Trends Survey on behalf of Masergy

<sup>2</sup>Vertical Systems Group STATFlash: 'How will COVID-19 impact SD-WAN?' See <https://www.verticalsystems.com/2020/05/28/statflash-sdwan-covid-2020/>

Enterprises continue to show a strong preference for managed SD-WAN services, with a combined 75 percent of respondents to IDG's 2020 SD-WAN Market Trends Report indicating that they currently utilize fully managed or co-managed service.<sup>3</sup> In line with these responses, Vertical System Group estimates that the number of billable U.S. installations of carrier managed SD-WAN services saw an 89 percent increase in 2019 - and they also recognized Comcast as a leader in this market.<sup>4</sup>



**Figure 2 – SD-WAN management preferences, source IDG**

### 3. Managed SD-WAN Service Evolution

For their initial managed SD-WAN offers, CSPs often deployed a single vendor's solution, which typically consisted of SD-WAN controller software and one or more Virtual Network Functions (VNFs), all bundled as a single package on an appliance or as software for deployment on a server.

As the SD-WAN market matures and segments, it is becoming clear that no one solution can address the full range of customer requirements and preferences.<sup>5</sup> Furthermore, CSPs that rely on a single vendors SD-WAN solution face heightened business risks: what happens if their vendor is acquired, fails to keep pace with market evolution, lacks one or more critical features, or develop support or quality issues?

Unfortunately, broadening their managed SD-WAN service portfolio is operationally challenging. In addition to investing resources to qualify and select vendors, the new solutions must be evaluated, tested and integrated with the B/OSS, new service design process must be established, and new methods of procedure (MOP) must be created for service activation and assurance. Additionally, activating and assuring managed SD-WAN services using traditional operations is both inefficient and expensive.

<sup>3</sup> IDG Market Pulse Research 2020 SD-WAN Market Trends Report: 'New Normal' Puts SD-WAN Digital Transformation in the Fast Lane" available at [www.masergy.com/white-paper/2020-sd-wan-market-trends-report](http://www.masergy.com/white-paper/2020-sd-wan-market-trends-report)

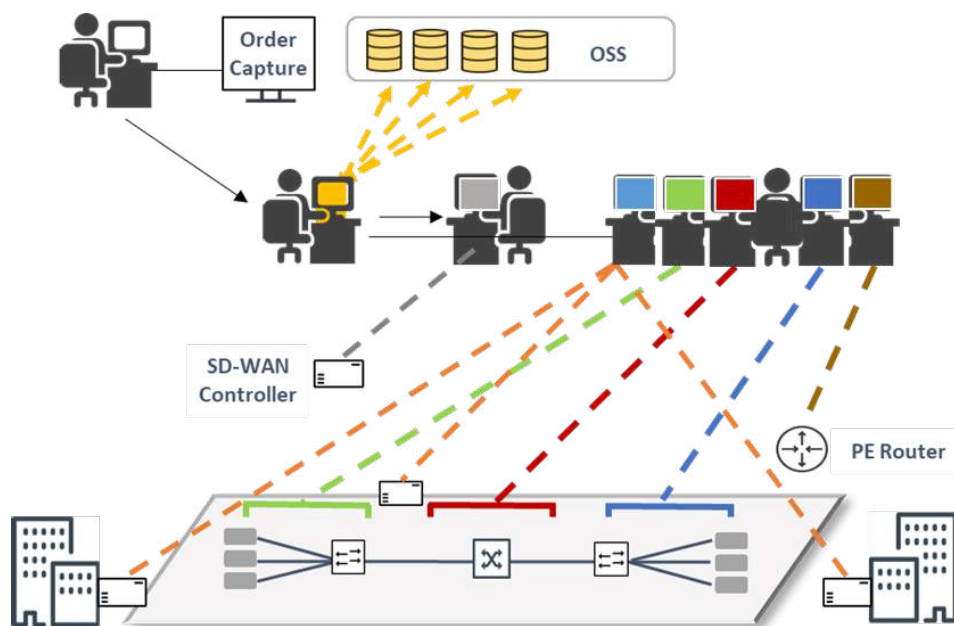
<sup>4</sup> Vertical Systems Group, 2019 U.S. Carrier Managed SD-WAN Services Leaderboard, see [www.verticalsystems.com/2020/04/21/2019-us-sd-wan-leaderboard/](http://www.verticalsystems.com/2020/04/21/2019-us-sd-wan-leaderboard/)

<sup>5</sup> WANSPEAK blog: 'Service Providers Need More Than One Option for Managed SD-WAN', see <https://blog.silver-peak.com/service-providers-need-more-than-one-option-for-managed-sd-wan>

## 4. Managed SD-WAN using Legacy Operations

When a customer order is processed, relevant portions of the order are forwarded to the engineering team so that they can design the service and create a related provisioning work-order. To do this, engineers must gather data from the B/OSS as well as Network and Element Management Systems (NMS / EMS). These manual processes take hours per service endpoint, and it can take weeks or months to complete these tasks for a customer with dozens or even hundreds of branch locations.

Once complete, the work-order is split into task-specific work orders, each with their own MOP, and forwarded to technicians that will use Command Line Interfaces (CLI's) or NMS / EMS to configure the SD-WAN controller, SD-WAN appliances or uCPEs, VNFs, PE routers, and inter-provider interfaces. This manual configuration is tedious and time-consuming - configuring the SD-WAN controller alone can involve hundreds of inputs - and the process must be repeated for each branch sites that is part of the service. Further complicating this process, there are dependencies for workflow execution: equipment must be shipped, technicians scheduled, access to facilities arranged. Once again, this process takes weeks or months to schedule and complete.



**Figure 3 – Manual SD-WAN service and network configuration**

Once the service is active, technicians must design and provision network and service monitoring and assurance. This is also a complex, manually intensive, and time-consuming activity, involving a wide variety of physical and virtual elements, network technologies, probes, and monitoring systems.

Because SD-WAN is an overlay technology that is disconnected from the underlay network, it is also difficult to use traditional trouble-to-resolve processes when performance issues occur. Fault isolation requires correlating data from a variety of independent service (overlay) and network (underlay) monitoring systems, making troubleshooting efforts slow, resource intensive, and frequently inaccurate. As a result, non-critical issues may become critical - and even service affecting - because they were not addressed correctly or in a timely manner.

## 5. The Business Impact of Legacy Operations

While Managed SD-WAN services are ideally matched to today's enterprise communications needs, designing, delivering, managing and assuring these services using traditional operational processes falls short of today's business environment, which calls for more velocity, agility, and visibility, at lower cost, from both the CSP and customer perspective.

Major limitations include

- Manual service design and work-order creation is slow, expensive, and prone to human error.
- Manual order process management is slow, opaque and error prone, leading to order fallout.
- Manual B/OSS interrogation and rekeying data into work-orders is slow, costly, and error-prone.
- Manual network configuration and service provisioning are slow, expensive, and error-prone—especially for NFV-based services—resulting in multi-week SD-WAN service activation per site.
- Manual intervention in fault isolation and remediation processes – especially alarm and event correlation between the SD-WAN overlay and network underlay – is too slow and difficult.

These limitations and related issues add cost and delay to the SD-WAN service lifecycle, which negatively impacts margins, impedes revenue recognition, and return on investment (ROI), and diminishes customer satisfaction. Worse yet, these issues are duplicated when a CSP adds new SD-WAN solutions to their portfolio, further increasing costs, complexity, and the burden on technical personnel.

## 6. SD-WAN-related Standardization Efforts

The MEF and the TM Forum and other organizations have developed standards to help CSPs build, deliver and manage services quickly and cost-effectively, using open interfaces and repeatable process. In large part, these efforts are complimentary and converge on similar architectures, protocols, and interfaces.<sup>6</sup>

### 6.1. MEF

The MEF<sup>7</sup> is an industry association of 200+ member companies that collaborate to enable the automated delivery of standardized services within and across multiple CSPs. SD-WAN service standardization has been conducted as part of the MEF 3.0 Global Services Framework.<sup>8</sup>

Important SD-WAN related MEF standardization activities include:

- MEF 70<sup>9</sup> - SD-WAN Service Attributes and Services – is the industry's first SD-WAN standard. It defines Managed SD-WAN Services as a specific use case for a MEF Third Network service and describes the requirements for an application-aware, overlay WAN connectivity service that uses policies to determine how application flows are directed over multiple underlay networks.

---

<sup>6</sup> Note that current standards and standards efforts at these two organizations are not designed to permit interoperability between SD-WAN solutions from different vendors.

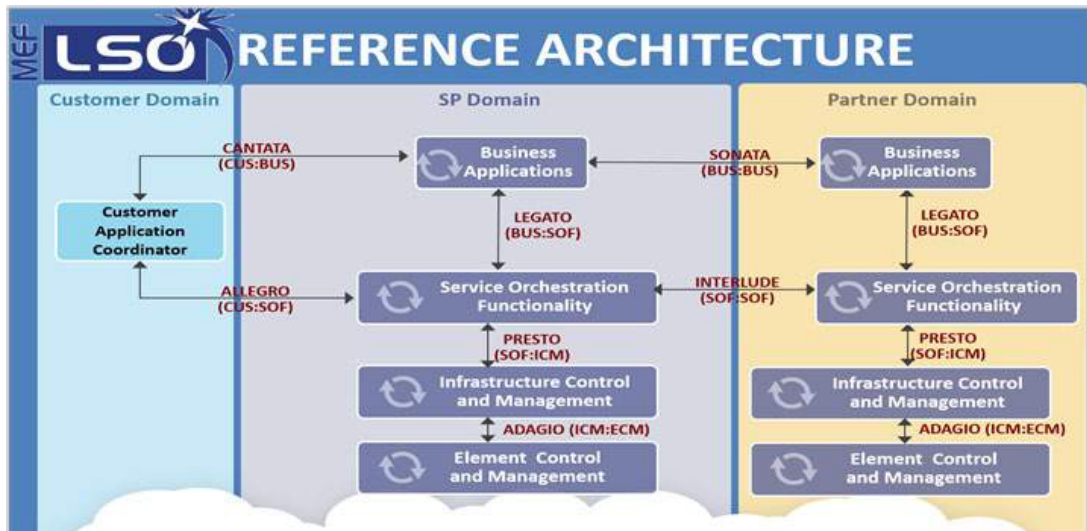
<sup>7</sup> To learn more about the MEF, see [www.mef.net/about-mef](http://www.mef.net/about-mef)

<sup>8</sup> MEF 3.0, see [www.mef.net/mef30/overview](http://www.mef.net/mef30/overview)

<sup>9</sup> MEF 70 at [www.mef.net/resources/technical-specifications/download?id=122&fileid=file1](http://www.mef.net/resources/technical-specifications/download?id=122&fileid=file1)



- MEF W90<sup>10</sup> - draft SD-WAN Certification Test Requirements – is part of the broader MEF 3.0 certification program. It tests the service attributes and their behaviors as defined in MEF 70.
- MEF 55<sup>11</sup> defines a Lifecycle Service Orchestration (LSO) framework for standard, automated service lifecycle orchestration that applies to all MEF 3.0 services, including SD-WAN.



**Figure 4 – MEF LSO Reference Architecture diagram**

MEF LSO proposes to automate the service lifecycle across all network domains responsible for delivering MEF 3.0 Network Connectivity Services, which includes SD-WAN. As figure 4 indicates, MEF LSO also defines Management Interface Reference Points between LSO functional management entities (e.g., Cantata, Allegro, Legato, etc.); application programmable interfaces (APIs) at these reference points to facilitate service orchestrations and automation.

## 6.2. TM Forum

The TM Forum<sup>12</sup> is an association of over 850-member companies' customers across 180 countries that drives collaborative problem-solving to help CSPs transform their business operations, IT systems and ecosystems. Their Open Digital Architecture (ODA)<sup>13</sup> provides a layered architectural approach that separates areas of concerns and uses standard REST-based Open APIs to expose services within a layer to adjacent layers. This approach is markedly different from the tight vertical integration between the network and the B/OSS that is implemented in traditional CSP operations environments.

<sup>10</sup> MEF W90 - SD-WAN Certification Test Requirements at [www.mef.net/mef-3-0-service-technology-certification](http://www.mef.net/mef-3-0-service-technology-certification).

<sup>11</sup> MEF 55 at [www.mef.net/Assets/Technical\\_Specifications/PDF/MEF\\_55.pdf](http://www.mef.net/Assets/Technical_Specifications/PDF/MEF_55.pdf)

<sup>12</sup> To learn more about the TM Forum, see [www.tmforum.org/about-tm-forum/](http://www.tmforum.org/about-tm-forum/)

<sup>13</sup> See [www.tmforum.org/oda/](http://www.tmforum.org/oda/)



**Figure 5 – TM Forum Open Digital Architecture framework**

TM Forum’s suite of 50+ REST-based Open APIs are reusable and enable standardized communications with B/OSS and management systems, making service creation and ongoing operations simpler and more efficient; additionally, they can be used to expose capabilities to partners. The use of standardized APIs also reduces the costs and risks associated with traditional B/OSS integration using proprietary interfaces.

As of July 2019, 18 leading CSPs and 48 leading technology vendors have signed the Open API Manifesto,<sup>14</sup> committing to include TM Forum Open APIs in their RFPs, and in products, respectively.

## 7. Automating SD-WAN operations

Today, CSPs can leverage standards from the MEF and TM Forum, together with advances in SDN and NFV technology to onboard new SD-WAN solutions and streamline the design, activation, and assurance of managed SD-WAN services in a closed-loop.

### Key Components of Closed-Loop SD-WAN Service Lifecycle Automation

The following systems work together to provide SD-WAN service lifecycle automation in a closed-loop.

- **Service order management system:** communicates with the Order Management System (OMS) and customer portals; maintains the service catalog; decomposes work orders; communicates with the orchestration system for order execution; coordinates the order workflow; and provides visibility into workflow status and issues, among other tasks and responsibilities,
- **Orchestration system:** communicates with a service order management system or OMS to receive service activation orders; orchestrates service design; communicates with SD-WAN controllers, domain layer controllers and NMS / EMS to configure provider edge (PE) routers, uCPE and other

<sup>14</sup> See [www.tmforum.org/open-apis/open-api-manifesto/](http://www.tmforum.org/open-apis/open-api-manifesto/)

elements; instantiate and chain VNFs as well as manage the VNF lifecycle; and communicate with assurance and analytics systems in a closed-loop, so that it can dynamically modify network paths to optimize resources or respond to issues when they arise.

- **Assurance and analytics system:** aggregates and analyzes network telemetry; monitors network and service health; dynamically correlates alarms and events from the SD-WAN overlay and network underlay to identify root-cause issues; and communicates with the orchestration system to enable resource optimization and to identify issues when they arise.

In addition to the basic functions listed above, these components must support open APIs, and provide a DevOps style ability to add new open APIs quickly and easily. Ideally, these components should also provide comprehensive visualization tools – graphically displaying workflow status, service path mapping across heterogeneous vendors, elements, and domains, offering a single unified view for monitoring all the resources.

Of course, beyond functionality, CSPs must ensure their vendors have the proven ability to support their systems and solutions at scale—large enterprises can have branches around the globe, and high expectations of pre- and post-deployment support.

Many different vendors provide one or more components that enable closed loop SD-WAN service lifecycle automation, and TM Forum Open APIs enable communications between these systems. This approach permits vendor innovation, and avoids vendor lock-in, and lets CSPs and their systems integration partners build ‘best of breed’ solutions that are closely aligned with their business goals.

CSPs should also be able to learn more about vendor capabilities by researching public proof of concept demonstrations, such as TM Forum Catalysts.<sup>15</sup> Another worthwhile and simple area to investigate is standards-compliance certification programs, such as those maintained by the MEF<sup>16</sup> and TM Forum<sup>17</sup>.

## 8. Standards-based Managed SD-WAN Automation

Figure 6 depicts how a CSP can implement a MEF LSO architecture and use standardized APIs as the foundation for automated SD-WAN operations. MEF LSO management reference points are defined North – South between layers (e.g., between the business layer and service orchestration layer; and between the service orchestration layer and the infrastructure control and management layer, which includes the SD-WAN controller and domain controllers).

A variety of APIs can be used between different layers, including TMF APIs (TMF 641 Service Ordering API is depicted between business and service orchestration layers; the Open Networking Foundations (ONF) Transport API (TAPI) and several others are depicted between the service orchestration layer and the infrastructure control and management layer); TMF APIs are also depicted between systems within the service orchestration layer (TMF 641 and TMF 623 SLA Management API are depicted).

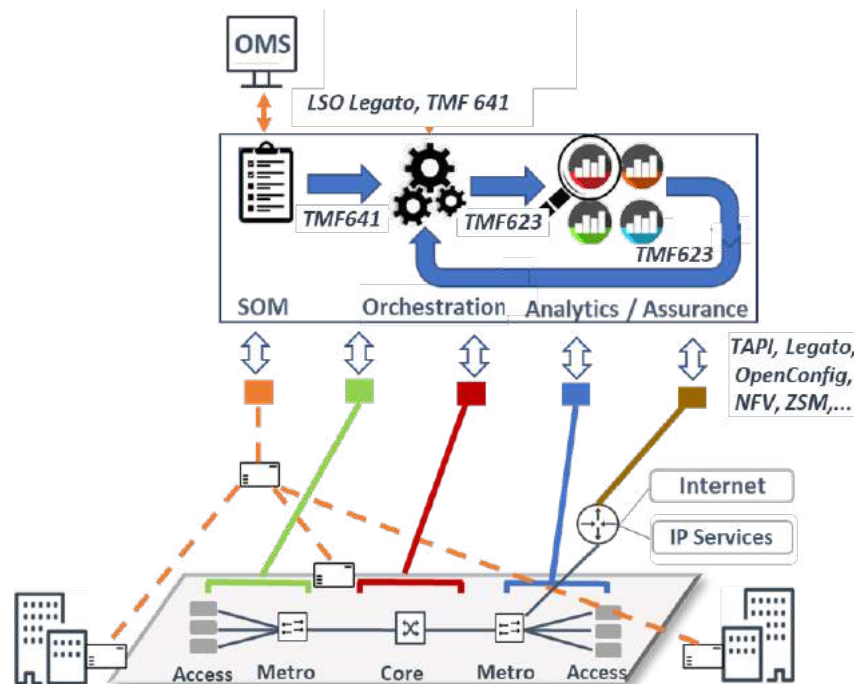
---

<sup>15</sup> TM Forum Catalyst program at [www.tmforum.org/collaboration/catalyst-program/home/](http://www.tmforum.org/collaboration/catalyst-program/home/)

<sup>16</sup> MEF Certification program at <https://www.mef.net/certification/mef-certification-programs>

<sup>17</sup> TM Forum conformance overview at <https://www.tmforum.org/conformance-certification/open-api-conformance/> and vendor leaderboard at <https://www.tmforum.org/conformance-certification/open-api-conformance/#table>

In this simplified example, the service order management system receives an order from the OMS and communicates with the orchestration system for service design. It then decomposes the work order and coordinates the order workflow execution. The orchestration system sends configuration and provisioning instructions to the SD-WAN controllers and domain layer controllers in the infrastructure control and management layer for execution.



**Figure 6 – SD-WAN Automation framework**

Once the service is active, assurance is dynamically instantiated end-to-end across the service path. Remediation for service affecting conditions is typically triggered dynamically, while analytics automates the evaluation of all alarms and events to identify root cause conditions. Issue resolution can be fully automated, or a mix of human / automated operations can be used, based on policy.

## 9. The Business Impact of Standardized, Automated Operations

When CSPs implement standardized, automated operations to design, deliver and assure their managed SD-WAN services, there are significant advantages for both the CSP and their customers.

From the CSP perspective, the major advantages include:

- Standard-based APIs eliminate expensive proprietary B/OSS integration activities, catalog and APIs are re-used across use-cases, avoiding wasteful one-off development efforts.
- Automated order process management that eliminates hundreds of tedious, manual, and error-prone tasks and avoids expensive design errors that can negatively impact customer commits.
- Automated orchestration accelerates network configuration and service provisioning, especially for NFV-based and hybrid services—enabling rapid service activation per site.
- Alarm and event correlation between the SD-WAN overlay and network underlay streamlines troubleshooting and fault isolation activities and helps protect service level agreements (SLAs).

These advantages reduce cost and delay throughout the entire SD-WAN service lifecycle, helping CSPs improve margins, accelerate time to revenue and ROI, and increase customer satisfaction. Additionally, APIs, templates and MOPs can be reused with any new SD-WAN vendor solution, which accelerates new vendor and solution introductions.

## 10. Managed SD-WAN Automation in Action

Let us look at a few CSPs and recent catalyst demonstrations that have extensively incorporated TM Forum and MEF standards, and automation, into their managed SD-WAN services offerings

### TMF Catalyst Demonstration: Automating Network as a Service

SD-WAN was featured in the *Automating Network as a Service using Operational Domain Management (ODM), TM Forum Open APIs and MEF LSO Catalyst*,<sup>18</sup> which was championed by Telstra, Vodafone, Orange Business Services, AT&T and PCCW. The demo included DGIT Systems order management solution and the Blue Planet orchestration platform - both supported by Infosys – and leveraged MEF LSO with TM Forum Open APIs to dynamically activate SD-WAN and other MEF 3.0 services.

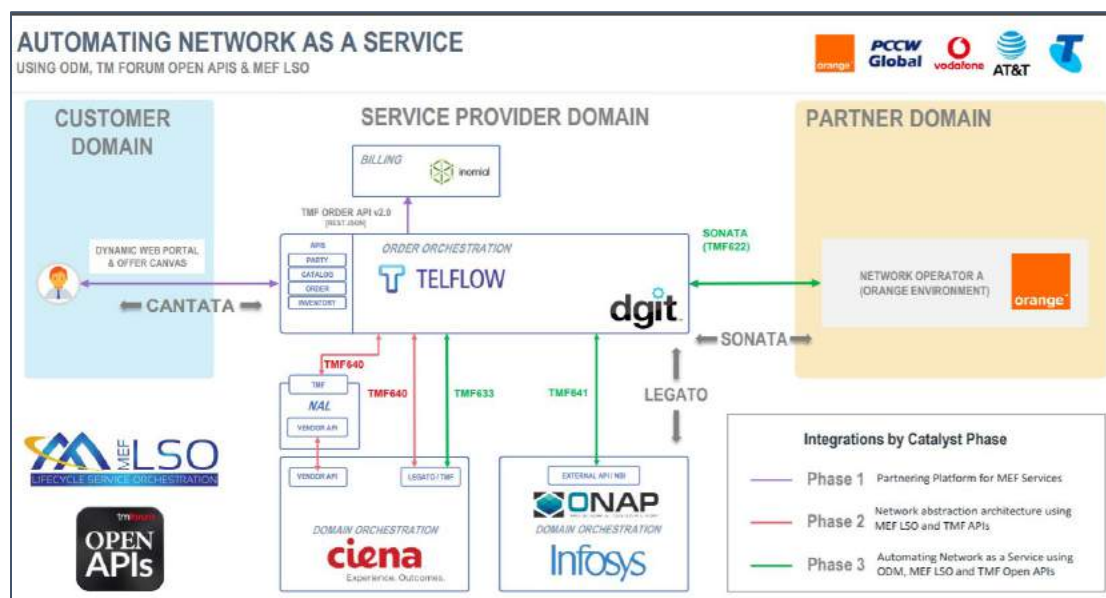


Figure 7 – Automating Network as a Service Catalyst reference diagram

<sup>18</sup> See <https://www.tmforum.org/wp-content/uploads/2017/02/T.-Automating-Network-as-a-Service.pdf>

## Comcast Business

Comcast Business - a subsidiary of Comcast Corporation, the second-largest broadcasting and cable television company in the world by revenue - was among the first cable operators to launch an SD-WAN service, among the few to achieve MEF 3.0 SD-WAN service certification, and the only MEF 3.0 SD-WAN Services certified CSP on the 2019 U.S. Carrier Managed SD-WAN Services Leaderboard.<sup>19</sup>

Their managed SD-WAN service is built on the ActiveCore SDN platform and leverages the internet or Comcast's nation-wide IP network as the network underlay, offering up to gigabit connectivity speeds.

Comcast Business is a 2020 LightReading Leading Lights Finalist in the Most Innovative SD-WAN Service category, and recently announced they would integrate Fortinet's virtual firewall appliances with ActiveCore to improve SD-WAN security by addressing threats such as malware.

## Windstream Enterprise

Windstream Enterprise is a leading provider of advanced network communications and technology solutions to enterprise customers across the U.S. and ranks fourth on Vertical Systems Group's 2019 U.S. Carrier Managed SD-WAN Services leaderboard.

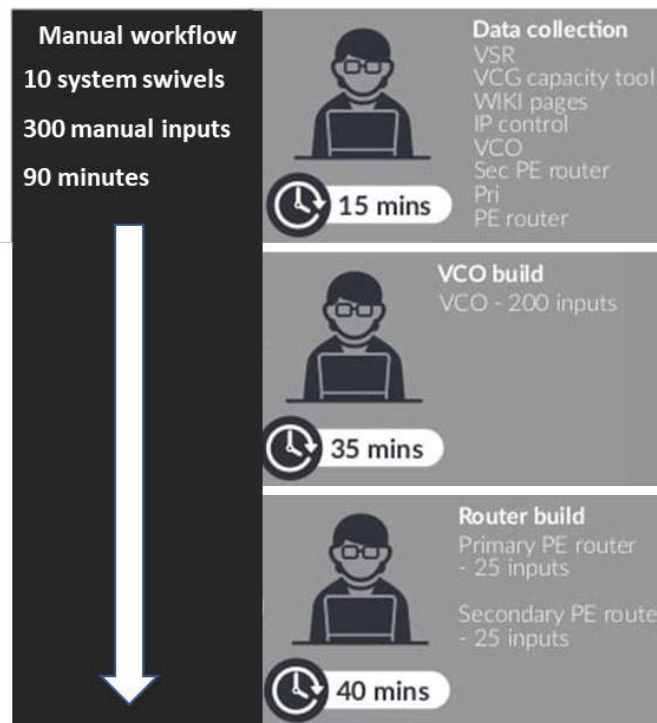
As highlighted in a TM Forum case study,<sup>20</sup> Windstream's initial SD-WAN deployment relied on a manual service activation process, which had technicians collecting data from 10 independent systems, and then rekeying that data into the SD-WAN controller and PE router to configure those devices.

---

<sup>19</sup> See [www.verticalsystems.com/2020/04/21/2019-us-sd-wan-leaderboard/](http://www.verticalsystems.com/2020/04/21/2019-us-sd-wan-leaderboard/)


<sup>20</sup> See <https://inform.tmforum.org/casestudy/windstream-use-intelligent-automation-to-cut-provisioning-time-by-80/>





**Figure 8 – Windstream Enterprise legacy operational workflow for SD-WAN Services**

Recognizing that their outdated operational processes were constraining their market success, Windstream implemented TM Forum ODA and related APIs and an orchestration platform that automated service activation. Ultimately, Windstream reduced systems accessed by 90%, manual inputs by 99.7%, and the total number of technician minutes required for service fulfillment by 98.9%.



**Figure 9 - Windstream Enterprise's automated SD-WAN service activation workflow**

## 11. Conclusion

The SD-WAN market has moved into the mainstream, with even large conservative enterprises adopting managed SD-WAN services. Analysts are forecasting high double-digit growth rates through at least 2024, and while the current global pandemic may disrupt growth in the near-term, it is likely to accelerate adoption in the long term.

To stay ahead of the competition, CSPs must move beyond basic SD-WAN services and single-vendor solutions. Fusion Connect, Windstream and GTT Communications<sup>21</sup> are just three public examples of CSPs with multi-vendor managed SD-WAN portfolios, there are many others – including some with more than two vendors.

Of course, a broad portfolio alone is not enough. CSPs must also differentiate by providing a better customer experience through rapid delivery, deeper customer insight, and superior availability. And they must provide this experience efficiently if they are to fully profit from their efforts and investments.

Achieving a broad and compelling managed SD-WAN service portfolio and a superior, streamlined operations environment requires CSPs to implement open standards and automation – especially closed-loop service lifecycle automation. This modern approach allows CSPs to add new vendors and services without penalty, and removes slow, expensive, and error-prone manual intervention from SD-WAN service design, activation, and assurance processes.

Based on the progress that has been made at the MEF and TM Forum - with significant contributions from cable operators - CSPs can now deploy advanced orchestration, assurance and analytics systems in an open, modular and standards-based automation framework that addresses the entire SD-WAN service lifecycle.

CSPs that implement such solutions will be able to reduce operational expenses, maximize resource utilization and increase staff efficiency even as they add multiple SD-WAN vendors to their portfolio. Furthermore, they will be able to accelerate time to revenue and improve customer satisfaction, enabling them to lead and thrive in competitive markets, and adapt to ever-evolving market requirements.

---

<sup>21</sup> Fusion Connect recently announced the expansion of their managed SD-WAN service portfolio beyond VMware VeloCloud; they will soon offer the Fortinet Secure SD-WAN solution, which features a built-in next-gen firewall, web filtering and an intrusion prevention system. See <https://www.channelpartneronline.com/2020/08/26/fusion-connect-picks-fortinet-as-second-sd-wan-partner/>



# Abbreviations

|        |                                               |
|--------|-----------------------------------------------|
| APIs   | application programmable interfaces           |
| B/OSS  | business / operations support system          |
| CLI    | command line interface                        |
| CPE    | customer premise equipment                    |
| CSP    | communications service provider               |
| IP     | Internet protocol                             |
| ISBE   | International Society of Broadband Experts    |
| LSO    | lifecycle service orchestration               |
| MOP    | method of procedure                           |
| MSO    | multi services operator                       |
| NCTA   | The Internet & Television Association         |
| ODA    | Open Digital Architecture                     |
| OMS    | Order management system                       |
| ONF    | Open Networking Foundations                   |
| OPEX   | operational expense                           |
| PE     | provider edge                                 |
| REST   | representational state transfer               |
| ROI    | return on investment                          |
| SaaS   | software-as-a-service                         |
| SCTE   | Society of Cable Telecommunications Engineers |
| SDN    | software defined network                      |
| SD-WAN | Software-Defined Wide Area Network            |
| SLA    | service level agreement                       |
| SOM    | service order management                      |
| TAPI   | transport API                                 |
| uCPE   | universal customer premise equipment          |
| VNF    | virtual network function                      |

# Bibliography & References

## MEF

- MEF 55 - LSO Reference Architecture at [www.mef.net/Assets/Technical\\_Specifications/PDF/MEF\\_55.pdf](http://www.mef.net/Assets/Technical_Specifications/PDF/MEF_55.pdf)
- MEF 70 - SD-WAN Service Attributes and Services at [www.mef.net/resources/technical-specifications/download](http://www.mef.net/resources/technical-specifications/download)
- MEF W90 - SD-WAN Certification Test Requirements at [www.mef.net/mef-3-0-service-technology-certification](http://www.mef.net/mef-3-0-service-technology-certification)
- MEF YouTube Channel at [www.youtube.com/user/MEFCarrierEthernet](http://www.youtube.com/user/MEFCarrierEthernet)
- MEF 3.0 SD-WAN services page at [www.mef.net/mef30/overview](http://www.mef.net/mef30/overview)

## TM Forum

- TMF website at [www.tmforum.org/about-tm-forum/](http://www.tmforum.org/about-tm-forum/)
- Open Digital Architecture at [www.tmforum.org/oda/](http://www.tmforum.org/oda/)
- Open API Suite at [www.tmforum.org/open-apis/](http://www.tmforum.org/open-apis/)
- TM Forum YouTube Channel at [www.youtube.com/channel/UCLKFQ99UR0KRtF3BTQzurOw](http://www.youtube.com/channel/UCLKFQ99UR0KRtF3BTQzurOw)

- Windstream Case Study at <https://inform.tmforum.org/casestudy/windstream-uses-intelligent-automation-to-cut-provisioning-time-by-80/>

# **Accelerating the Virtualization:**

## **Introducing Hybrid Fiber Shelf into the Mix**

A Technical Paper prepared for SCTE•ISBE by

**Venk Mutalik**

Executive Director, HFC Architecture  
Comcast  
1401 Wynkoop St, Denver, CO  
+1 (860)-262-4479  
Venk\_Mutalik@Comcast.com

**Bob Gaydos, Comcast**

1800 Arch St, Philadelphia, PA  
Robert\_Gaydos@Comcast.com

**Dan Rice, Comcast**

1401 Wynkoop St, Denver, CO  
Daniel\_Rice4@Comcast.com

**Jorge Salinger, Comcast**

1401 Wynkoop St, Denver, CO  
Jorge\_Salinger@Comcast.com

# Table of Contents

| <b>Title</b>                                       | <b>Page Number</b> |
|----------------------------------------------------|--------------------|
| 1. Abstract .....                                  | 3                  |
| 2. Introduction to Virtualization .....            | 3                  |
| 3. Architectures and Virtualization.....           | 6                  |
| 4. Converging the Headend and the Fiber Plant..... | 10                 |
| 5. A Word about Coherent Optics .....              | 15                 |
| 6. Hybrid Fiber Shelf.....                         | 16                 |
| 7. Critical Infrastructure .....                   | 19                 |
| 8. Conclusions.....                                | 21                 |
| 9. Acknowledgements .....                          | 22                 |
| Abbreviations .....                                | 22                 |
| Bibliography & References.....                     | 23                 |

## List of Figures

| <b>Title</b>                                                                           | <b>Page Number</b> |
|----------------------------------------------------------------------------------------|--------------------|
| Figure 1 – Simplified Connectivity Diagram. Red=PHEs, Blue=SHes, N = Fiber Nodes ..... | 4                  |
| Figure 2 – End to end DAA Architecture Diagram .....                                   | 5                  |
| Figure 3 – HFC Architecture Diagram .....                                              | 6                  |
| Figure 4 – RFS Architecture Diagram.....                                               | 7                  |
| Figure 5 – HFS Architecture Diagram.....                                               | 7                  |
| Figure 6 – An example of Traditional CMTS and PHE to SHE to Node connectivity .....    | 8                  |
| Figure 7 - An example of Traditional vCMTS and PHE to SHE to Node connectivity .....   | 9                  |
| Figure 8 – Analyzing ~100K SHE-Node Fiber Links Across Our Footprint .....             | 10                 |
| Figure 9 – Converged Fibers .....                                                      | 11                 |
| Figure 10 – Taxonomy of the Optical Effects and Non-linearities.....                   | 12                 |
| Figure 11 – Full Spectrum Wavelength Planning .....                                    | 13                 |
| Figure 12 – Illustrating Analog and Digital Coexistence on the Same Fiber .....        | 14                 |
| Figure 13 – Coherent Optics Options .....                                              | 15                 |
| Figure 14 – Logical Block Diagram for HFS.....                                         | 17                 |
| Figure 15 – Measured RPD and EML Transmitter MER .....                                 | 18                 |
| Figure 16 – Node and RF Cascades .....                                                 | 19                 |
| Figure 17 – Summary of HFS .....                                                       | 19                 |
| Figure 18 – HFS Critical Infrastructure .....                                          | 20                 |

## 1. Abstract

Broadband access networks continue to experience customer growth and higher capacity demands year over year. Early and ongoing analysis of these demand trends enabled Comcast to develop a Distributed Access Architecture (DAA) that enabled us to keep up with capacity needs and enhance customer satisfaction with appreciable capital economics.

Successfully virtualizing complex portions of our network, including the CMTS, resulted in network simplification and the harmonization of multiple purpose-built platforms into one common entity. However, DAA deployments require headend and field modifications as part of the deployment. Even before the 2020 spike in capacity demands created by the COVID pandemic, and its consequent work from home requirements, efforts were underway to accelerate the virtualization, and associated benefits, of our networks, and to separate headend-centric innovations from those in field construction. These parallel efforts enabled us to press forward more rapidly in the adoption of virtualized CMTS technology.

A new concept was developed called a Hybrid Fiber Shelf (HFS) that integrates with virtualized CMTS/DAA Switches at one end, and with transmit/receive analog optical signals on the other end into the outside plant. This innovation provides rapid and economically sustainable increases in capacity, by independently accelerating vCMTS integrations while DAA construction proceeds in other areas of the network. HFS improves critical infrastructure in headends by saving wasted space, power and time inherent in RF combining and splitting circuits. Extending HFS into secondary headends (SHEs) also enables significant fiber reclamation. Locating these assets in secondary facilities closer to the traditional HFC nodes improves performance in ways that translate into enhanced capacity with upstream profile management tools also presented at this SCTE.

In this paper, we begin with a description of the optical and RF innovations that enabled a hybrid fiber shelf concept, and its theory of operations. We then describe the end performance and the impressive improvements in headend critical infrastructure. We finally describe the commonalities and nuances of this approach and its fit within the overall DAA architecture.


## 2. Introduction to Virtualization

A typical Cable System uses Cable Modem Termination Systems (CMTSs) to provide high availability upstream (US) and downstream (DS) signals between IP-connected devices (cable modems, gateways, set-tops) and IP-delivered services in digital Quadrature Amplitude Modulation (QAM) format. These QAM signals are RF-combined with other video sources, and then provided as inputs to high performance analog transmitters. Multiples of such analog transmitters' output wavelengths are then multiplexed into one fiber and sent to multiple Secondary headends. At the secondary headend, these signals are either demultiplexed and retransmitted, or amplified and distributed to individual nodes. After the node, multiple RF amplifiers in cascades amplify this signal to reach homes. In the reverse path, signals from the home arrive back at the RF amplifiers, which are bidirectional, and then reach the SHE from where they are often aggregated and sent back to the PHE to be terminated in the CMTS to complete the signal circuit.

The traditional CMTS is a formidable equipment type, with purpose-built hardware and significant scheduling and modulation software that has soaked over decades. All high-speed data (HSD) traffic flows through and terminates at the CMTS, making it a gatekeeper of broadband traffic for the MSOs.

Each time additional data was needed in the plant, typically in chunks of 50 MHz, the CMTS capacity would have to be augmented by combining additional ports, which creates capital costs.

This type of traffic infrastructure was fine while the HSD was a sliver of the total cumulative spectra delivered to the home, which was dominated by MPEG video, pay-per-view and VOD and fit the overall criteria of “pay as you grow” dictum of the MSOs. Readers may remember that long span of time when two 6 MHz channels served most of the voice and Internet services for most MSOs. However, when HSD usage started to grow exponentially, with CAGRs of ~35% year over year, and the CMTS ports started piling up faster and faster, the model of a traditional CMTS started becoming unwieldy. Something had to change, when HSD traffic dominated the US and DS spectra.

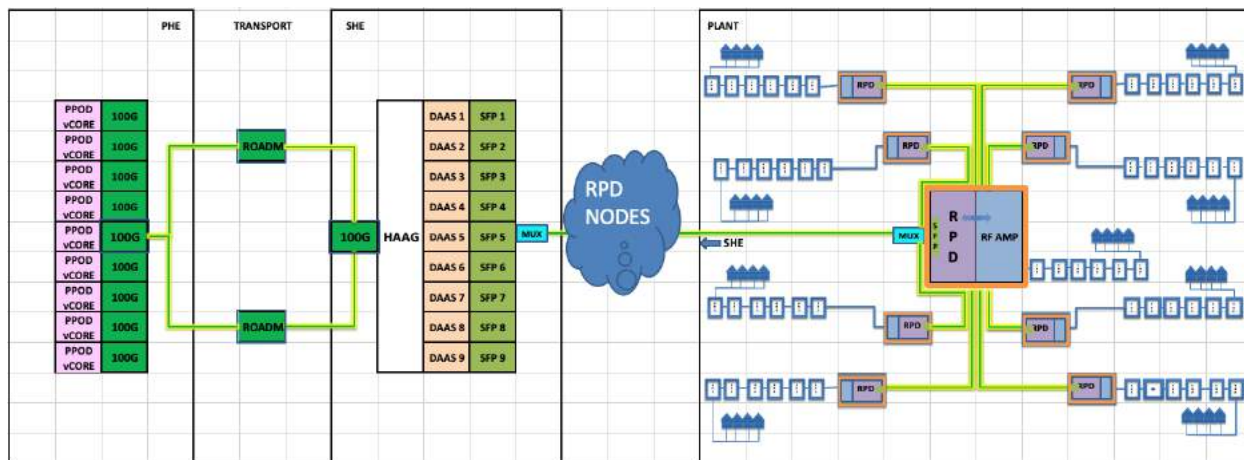


**Figure 1 – Simplified Connectivity Diagram. Red=PHEs, Blue=SHEs, N = Fiber Nodes**

It was around this time when we began to look at a new paradigm of network architecture. Distributed Access Architecture (DAA) envisioned the disaggregation of the CMTS into its component scheduling (aka MAC) and modulation (aka PHY) parts, accomplished by placing the PHY parts within the optical fiber node. Doing so would enable the headend to only concentrate on scheduling issues, while the PHY layer would concentrate on the modulation issues and thus distribute computational resources across the network, improve performance and reduce overall cost. The devices inside the nodes were called Remote PHY devices (RPDs) and thus was born the term “RPD node.”

For the RPD nodes to function well, the CMTS functionality in the headend must be allowed to scale for traffic from day one. Because the RPD has just one 10G optical connection, we would ideally like to have all traffic flow to it from one source, without any RF combining of other formats that would mitigate the benefits afforded by the RPD.

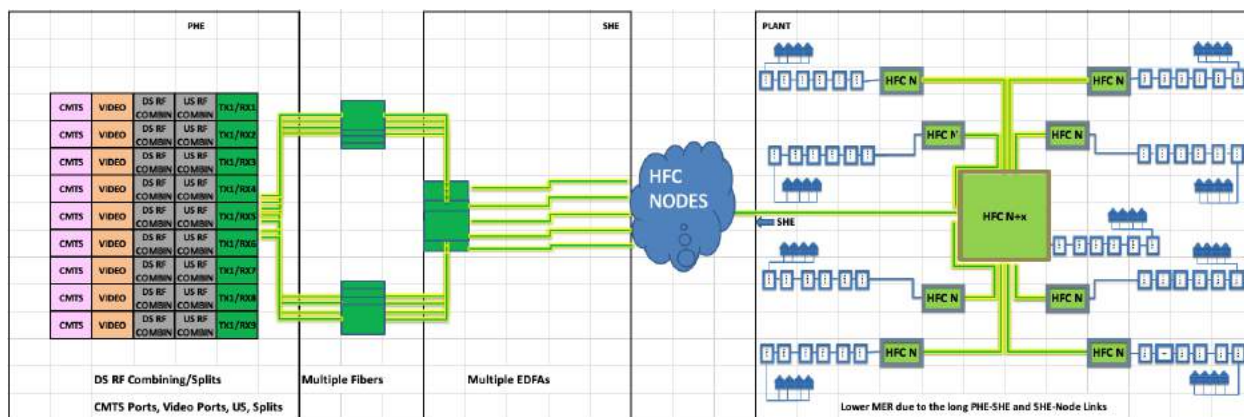
This is what prompted us to rethink the CMTS as a virtual machine: One that depends less upon purpose-built hardware that controlled the RPD, and thru it the Cable Modems (CMs) in the house. By creative use of software structures, such as Containers/Kubernetes, the entire CMTS MAC was thus virtualized and distributed across the network. The success of this effort is of breathtaking importance, because for the first time it allows MSOs to increase traffic on the networks without the additional CMTS port license fees for each upgrade.



**Figure 2 – End to end DAA Architecture Diagram**

Once the CMTS is virtualized, the core CMTS functionality, known as the virtual core (vCORE) or interchangeably the virtual CMTS (vCMTS), is located in the primary headend (PHE), while multiple Secondary Headends host what are called DAA Switches (DAASs). These DAAS ports are connected to the vCOREs thru high speed, sometimes coherent links, while the DAASs themselves are connected to the RPDs via 10G DWDM links. The RPD node then converts the 10 Gbps baseband data into QAM format and then distributes it into the networks. At the home, the CM receives this data, and sends its US data to the RPD, which converts that into baseband data and sends them to the DAAS ports in 10Gbps format. The DASS then sends it back to the vCORE via the high-speed links, completing the circuit.

DAA is just one type of architecture in market in Comcast. In reality, we support more traditional HFC nodes than the DAA kind. Unfortunately, for the traditional nodes, the vCORE-RPD node connectivity architecture does not work, because there is not RPD in the picture. Therefore, some of the best benefits of the vCORE are not available to parts of our footprint.



**Figure 3 – HFC Architecture Diagram**

Fortunately, there are ways in which this can be accomplished and in this paper we will discuss the basic idea of unleashing the vCORE onto a vast majority of our plant. Some of the ancillary benefits that accrue because of the vCMTS involve the significant critical infrastructure savings in the PHEs and SHE. Additionally, there is a potential for performance improvements over the existing link. Finally, we discuss a common wavelength plan that enables us to converge coherent, analog and 10G wavelengths onto one single fiber.

### 3. Architectures and Virtualization

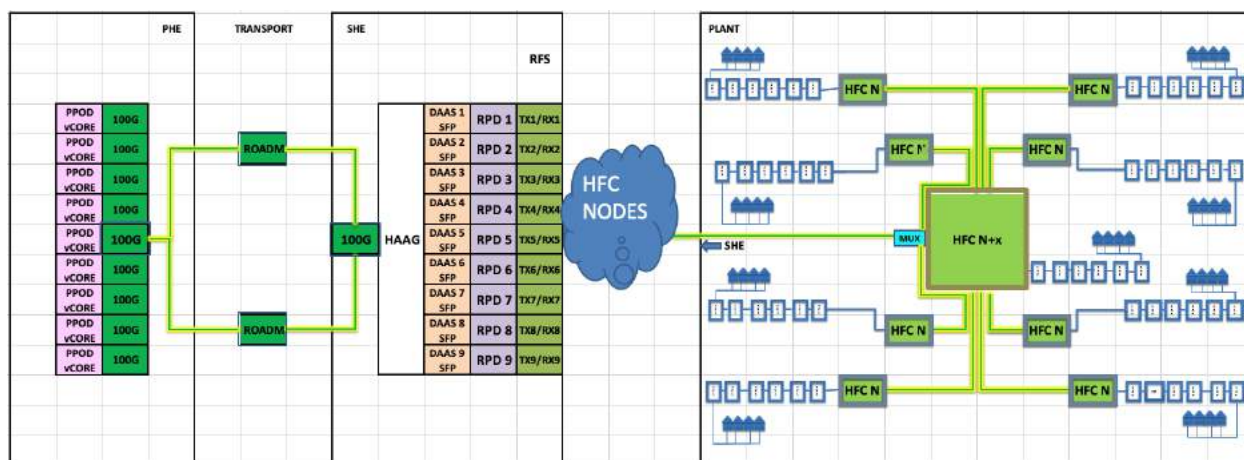
The key to the success of any standard is its ability to fit into and enhance multiple diverse architectures. Such is the case with virtualization. Virtualization of the cable modem termination system (CMTS) began as a way to enhance Distributed Access Architecture (DAA) strategies. However, as it turns out, virtualization is a great fit for multiple architectures. Accelerating the reach of virtualization enhance all these architectures.

We have already seen that in the DAA architecture, the RPD node is a perfect fit for vCORE, however we have a preponderance of the original HFC nodes. Therefore, it is essential to provide an option for using the vCMTS architecture, so as to realize its benefits in regular HFC nodes as well. It may well be asked why we would not go about and change the entire architecture to DAA; however, that is not such an easy task. A typical DAA architecture can also be a “node plus zero” (N+0) architecture, meaning no amplification/amplifiers after the node. A that goes into making a system conform to N+0, not the least of which is the ability to secure enough HHPs per RPD. While this might happen easily in high density areas, such is not the case for lower density areas. For those areas, Grey Optics Aggregation (GOA) architectures will come in handy to shore up the number of passings per RPD. However, in current environments, and especially with COVID resulting in a lot of work from home households, the capacity needs for both US and DS are quite substantial, and capacity must be improved on a timelier basis. This is one of the reasons to leave node placements and RF amplifier placements as is.

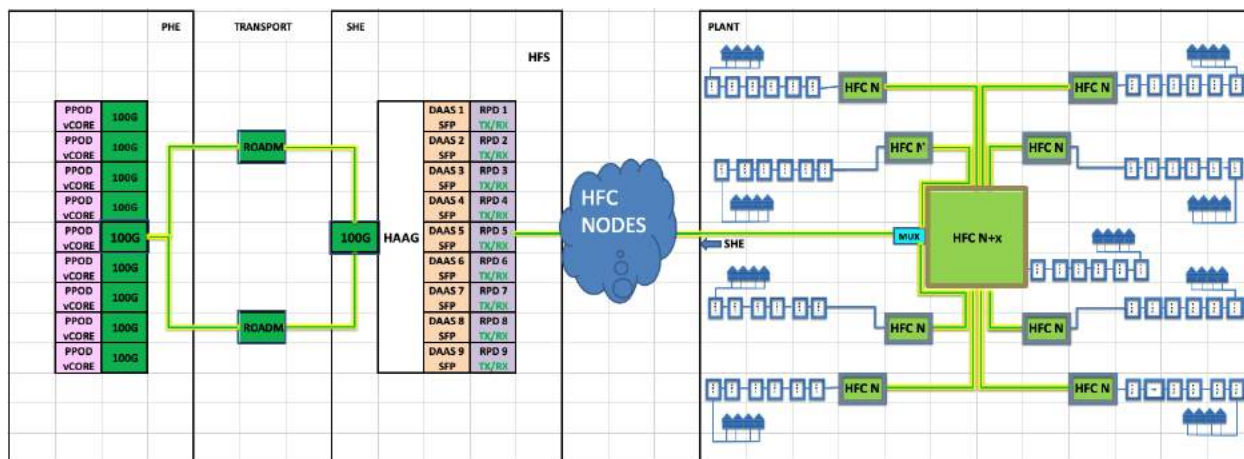
It may further be argued that at least the nodes could be replaced, leaving RF amplifier cascades alone. While that is a better option, even it will take time. RPD nodes are physically bigger and will almost certainly need a complete node re-splicing. As well, a training regimen will be needed for field technicians who may not yet be trained on RPD devices. In some cases, it will continue to make economic sense to leverage existing investments in analog nodes and optics, especially in cases where additional capacity can be leveraged. In any case it is quite hard and inconsistent to migrate RPDs in the field.



For these reasons, the best approach for some of the network segments is to move the RPD into the headend and co-locate these with the DAAS ports in the SHE. The RPDs are then connected directly to the DAAS ports with active optical cables (AOCs), and the output of the RPDs are then fed to DS analog transmitters (TXs). These TXs feed typical HFC analog optical nodes, which reach household CMs. The return signals from the CMs arrive at the node and get transported to the SHE, where they are fed to the return receivers, then connected to the RPDs, which are then fed back to the DAAS ports, completing the signal circuit. These returns from analog nodes can represent typical digital or analog signals.



**Figure 4 – RFS Architecture Diagram**

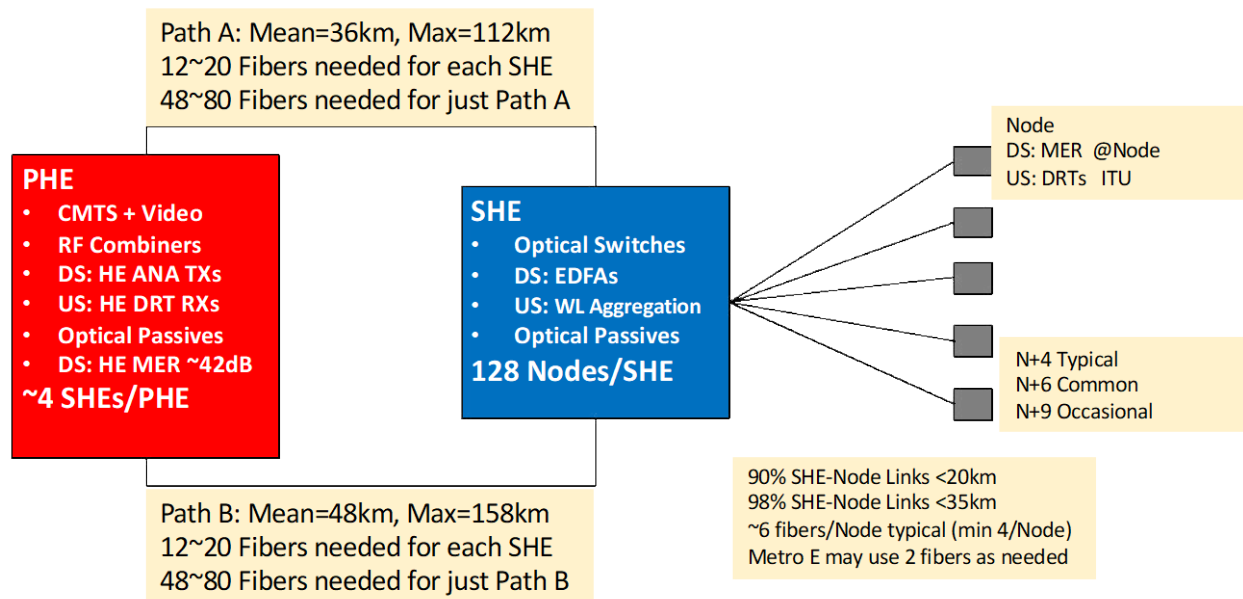


**Figure 5 – HFS Architecture Diagram**

This new approach to accelerate the rate of virtualization is called as the RPD Fiber Shelf (RFS) and Hybrid Fiber Shelf (HFS).

Figure 3 shows a set of details about a PHE and SHE within one of our regions. A general connectivity diagram for this was presented in Figure 1 Basic Connectivity Diagram. We see here that the PHE hosts the CMTS and video traffic, along with the associated RF combiners. This primary headend also houses the analog transmitters that are multiplexed and sent over redundant links to the secondary headend. At

the SHE, and opto-mechanical switch resolves the redundancy, and, after amplification, passes it over to the optical node.

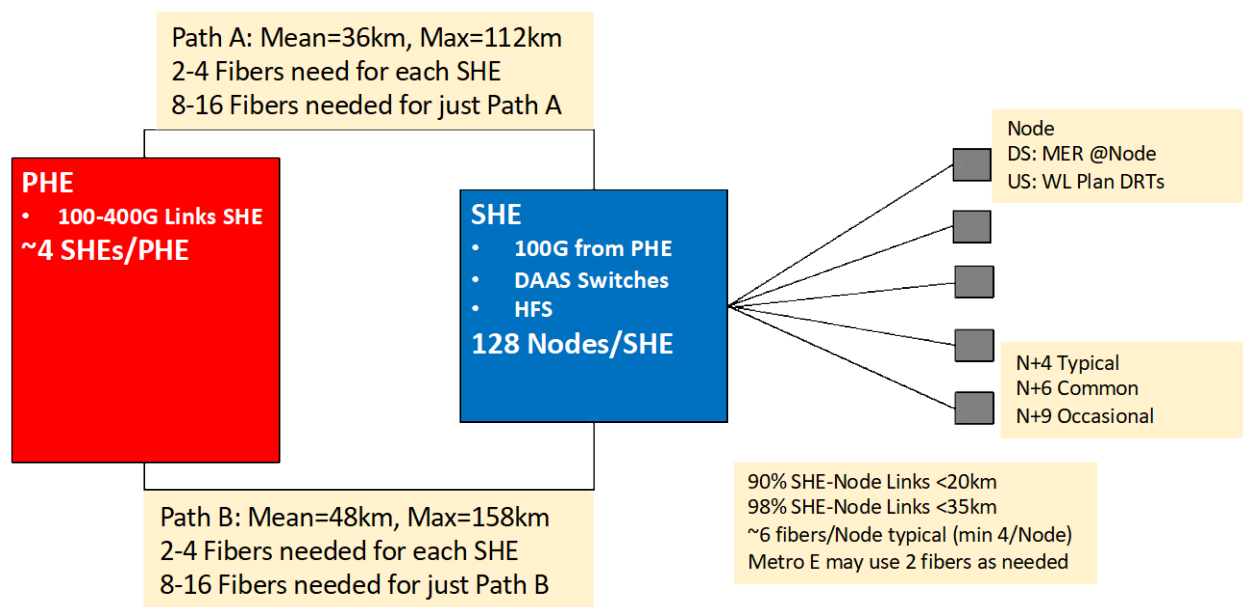


**Figure 6 – An example of Traditional CMTS and PHE to SHE to Node connectivity**

It must be noted here that due to the long reach from the PHE to the SHE and then to the neighborhood node, the Modulation Error Ration (MER) at the node itself is rather limited. Furthermore, the number of fibers need to traverse from the PHE to the SHE is substantial, because to avoid the deleterious fiber effects, analog transmission allows for 16 full spectrum wavelengths, which will be discussed in Section 4. In many cases, the rather large difference in primary and redundant routes creates a substantial MER delta when the redundant route becomes switched -- not to mention CMs generally reset themselves when there are substantial temporal changes caused by flight delay changes to the CMTS. The MER further reduces when we move from the node through the RF amplifiers.

In the return path, baseband digital returns have a rather large dynamic range, which gives them long reach capabilities. In this case, they are all combined together, and their light is shipped back to the primary headend. In any case, the number of fibers used from the PHE to the SHE are considerable, and any additions in the number of nodes to accommodate node splits, for example, could easily overwhelm the fiber infrastructure. Adding fibers to the PHE to SHE infrastructure is “easier said than done” due to cost of the large number of needed redundant links and also of the time associated with acquiring permits and the subsequent construction.

Fortunately, a move to HFS solves all these issues. In the diagram below, when HFS is applied, the PHE “only” has high speed coherent links of ~100-400G that traverse to the SHE over redundant routes. Since these are a well-established technology, and higher capacity, they require only a small fraction of PHE to SHE fibers. At the SHE, we connect up the coherent links to the DAAS ports. Typically, up to 48 DAAS ports can be lit with a 1RU DAA switch box. The DAAS ports are connected to the HFS and the analog link, then connected to the fiber node. Note here that the analog link traverses only the link from the SHE to the node.




**Figure 7 - An example of Traditional vCMTS and PHE to SHE to Node connectivity**

By way of example, comparing the two figures immediately above, we see that there is a net savings of 10-16 fibers per PHE to SHE hop, and since there are typically 4 SHEs per PHE, this results in 40-64 fibers saved per SHE cluster. There is an equal number of secondary fibers that can be saved, as well. Some of our bigger markets may have as many as 100 SHEs, in which case the numbers of fibers saved on the (very hard to get) HE-SHE link would be up to ~1500 to 3000 when the primary and redundant routes are counted!

It is this huge savings in PHE/SHE fibers, and the associated reduction in critical infrastructure pressure at the PHEs, that accounts for a push towards HFS. But for this to happen, the HFS should be able to fit into a secondary headend – and some of them can be small and overpopulated. Furthermore, the link lengths from the SHE to the nodes are short enough for the HFS to offer good performance, even accounting for any RF cascades after it. All of these points are discussed in this paper, but we begin with a survey of our fiber links.

The graph shown in Figure 5 is the result of a survey of >100K geographically dispersed fiber nodes within our U.S. footprint. The average link length between the secondary headend to the fiber node is ~10 km; 98% of nodes within 30 km. This survey spans 16 major markets and all of our regional divisions.



**Figure 8 – Analyzing ~100K SHE-Node Fiber Links Across Our Footprint**

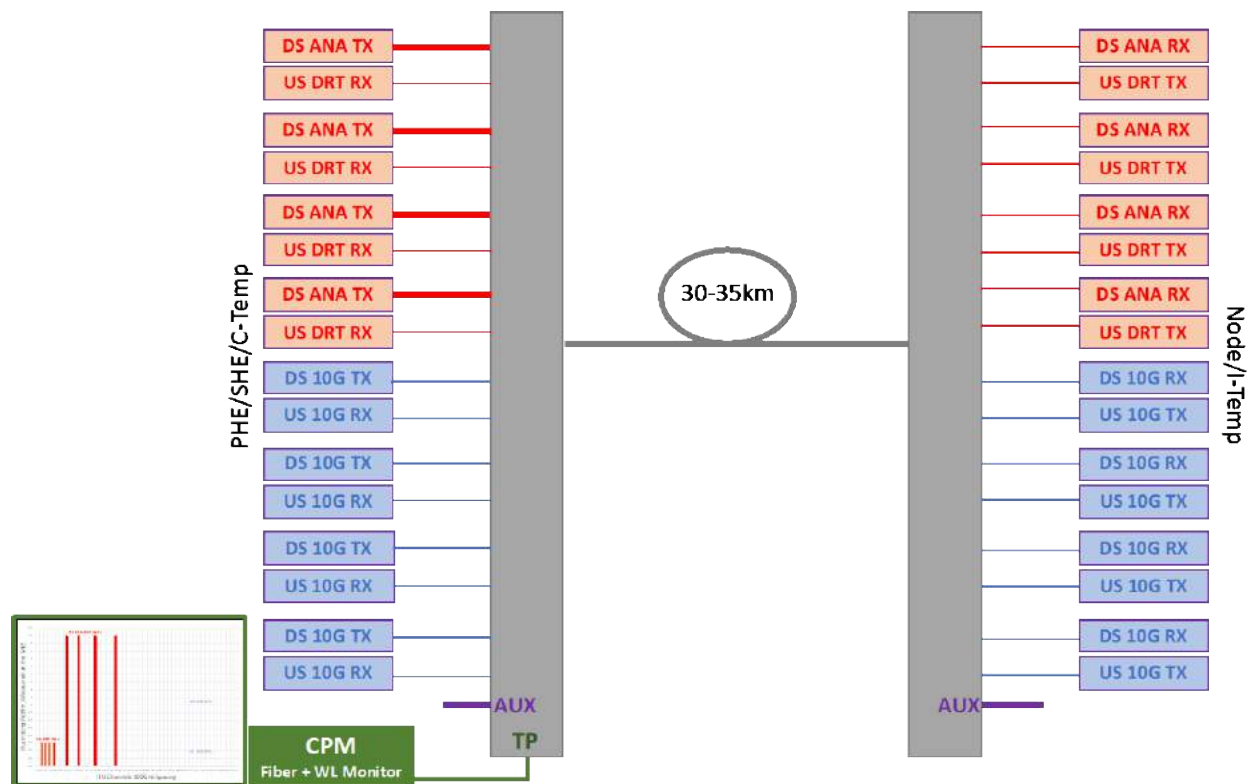
With this visual, it can be seen that the MER of the node for a signal generated at the SHE and fed with an analog transmitter could have a much higher MER than that of a long PHE to SHE link followed by a fiber run to the node, which will be shown in Section 6.

To recap, a move to the HFS will result in a significant reduction in fibers from the PHE to the SHE. With fiber lengths that are more modest, there also a substantial improvement in MER at the node. The adoption of a Hybrid Fiber Shelf also results in critical infrastructure benefits, including substantial powering and space reductions in secondary headends. These savings, combined with the speed with which HFSs can be deployed, and the resultant uniformity of vCMTS usage across the company, are perhaps its most attractive features.

## 4. Converging the Headend and the Fiber Plant

We have seen in earlier sections that the primary headend/PHE hosts vCOREs and coherent optics, and the secondary headend/SHE hosts the same high-speed optics, as well as the DAAS ports and analog optics. In many cases, the PHE also supports nodes directly, which are (aptly) known as “direct fed nodes.” In this case, the conventional CMTS is replaced with the vCORE; the DAASs are also connected to the vCORE and fed to the analog transmitters as described earlier. In this instance, all forms of optical conversions take place directly within one converged headend.

Similar to the case of a converged headend, we will encounter a substantial period of time when we might expect cases of converged fibers. By converged fibers, we mean fibers with bi-directional transmission of analog and digital wavelengths on the same single fiber. Consider the case shown in Figure 6, where a simplified diagram is presented.



**Figure 9 – Converged Fibers**

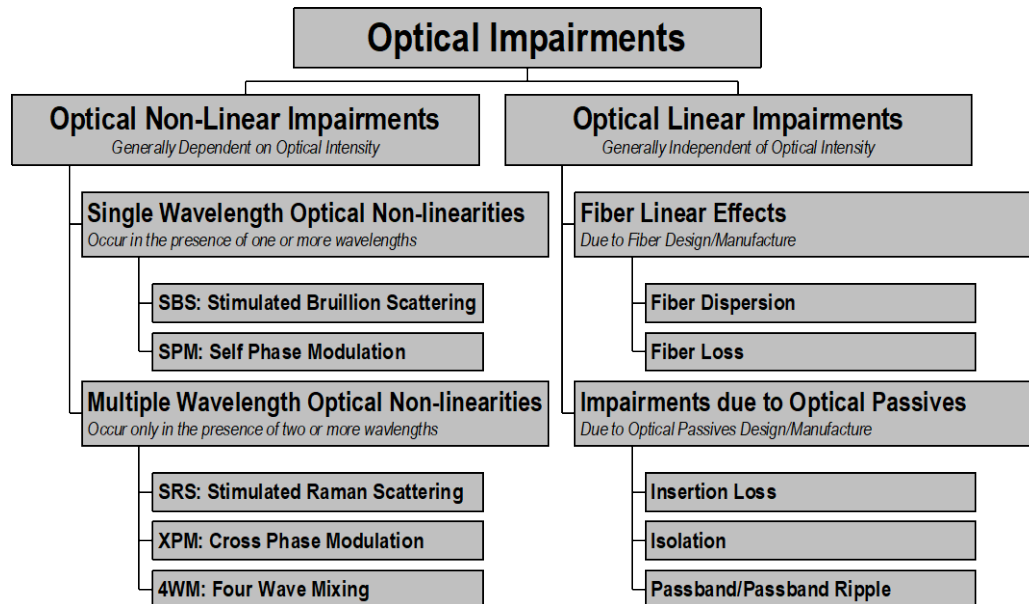
On the same fiber that hosts analog DS wavelengths, the US digital return transmitter (DRT) wavelengths may be hosted as well. Furthermore, 10G wavelengths that support DAA nodes and some other 10G links may need to support MetroE links and be hosted as well. For DAA operation, today we use 10G tunable SFPs.

Considering Figure 6 as part of the SHE, we can see that there are up to 4 DS analog transmitter wavelengths and up to 4 US DRT wavelengths, and up to 4 bidirectional 10G wavelengths circuits that could support any combination of DAA or MetroE circuits. A detailed description of the wavelength selection process is given in Figure 7. But here, we note that additional wavelengths that are part of a well thought out plan may still be added on the auxiliary port. In addition, the entire fiber and all sets of wavelengths are monitored over the consolidated test port by a continuous and pervasive monitor that is described in a different paper in the 2020 Cable-Tec Expo program.

There will be cases when over time, 100G/200G coherent wavelengths may need to be hosted on the same fiber, in order to enable Switch on a Pole (SOAP)-type architectures (this point will be discussed briefly in Section 5) In general, dual fiber options are not favored because the fiber coexistence use case presupposes a fiber scarcity that argues against this option. To enable multiple wavelengths of differing



power levels into the fiber, a comprehensive understanding of optical effects and non-linearities is needed.



**Figure 10 – Taxonomy of the Optical Effects and Non-linearities**

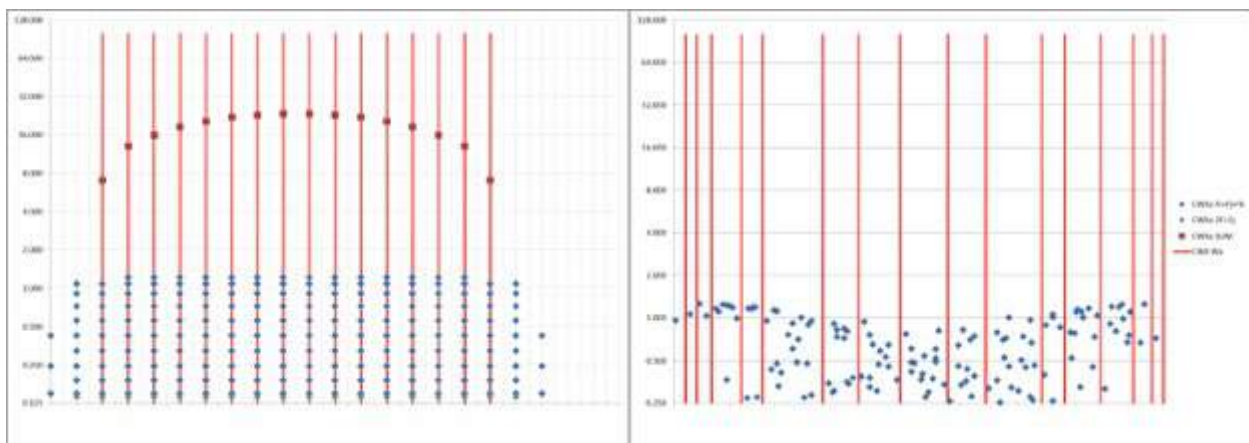
Presented in Figure 7 is a taxonomy of optical impairments. We begin by dividing impairments into linear (those that do not generally depend upon the optical power level) and non-linear (those that generally become worse with higher optical levels) impairments. Linear impairments can be attributable to the fiber itself, such as dispersion and fiber loss, which have to be dealt with in all of the signal types. For example, the use of externally modulated transmitters heavily reduces the effects of dispersion, since they are almost chirp free. For baseband signals, propagation length is inversely proportional to the square of the dispersion. Therefore, a 10G link that can go over 80km can only go over ~15km with 25G line rates, and less than a km at 100G with the same Non-Return-to-Zero (NRZ) signals. Transmitting the signals in the 1310 nm region, where the dispersion is close to zero, would work, but here we run into the other impairment -- the fiber loss -- which is considerably higher at 1310 nm than it is at 1550 nm, with no recourse for easy optical amplification.

Add to this the fact that optical passives also add to analog system impairments. Directly modulated lasers (DMLs) can be affected by passive ripples (undulations in the passband), however, externally modulated lasers do not have this effect because they are chirp free. For digital systems, the passband of the filters should be large enough and the adjacent channel isolation deep enough for successful transmission. We currently use thin film filters, as they are well suited to indoor and outdoor plant, and because we also have high and low wavelength counts distributed throughout the infrastructure.

A common theme here is that that for analog transmission, externally modulated transmitters seem to be a good fit. They are able to handle dispersion and optical passives better. In fact, well designed externally modulated lasers can have much better MER than DMLs could -- over longer reaches and over a wider variety of optical passives. The industry has used EMLs for quite a long time, but in recent years their prices have compared favorably to DMLs. As such they represent a great alternative for the HFS and provide appreciable performance benefits.

We now come to non-linear signal performance. These generally present as single wavelength impairments which will need to be optimized. For example, while EMLs behave well with dispersion and over the passives in the linear domain, they have very aggressive single wavelength non-linearities due to stimulated Brillouin scattering (SBS) which must be overcome. Fortunately for us, in HFS, we have modest fiber links, power levels are modest, and these have been overcome and are on par for the course. For digital transmission, these effects are minimal.

Finally, when we put multiple wavelengths on the same fiber, multiple nonlinear effects can take place simultaneously, such as Stimulated Raman Scattering (SRS), Cross Phase Modulation (XPM) and 4 Wave Mixing (4WM). SRS is mitigated by packing wavelengths closer; XPM is mitigated by moving wavelengths far apart; and 4WM is mitigated by avoiding all beat products that would produce significant SNR degradation by the effects of optical beat interference (OBI). OBI is a known issue in RF over Glass (RFoG) systems and has the ability to singlehandedly shut an entire system down.



**Figure 11 – Full Spectrum Wavelength Planning**


To make matters worse, 4WM is intermittent, seeing as how it depends upon beats lining up with carriers. Therefore a standard uniform wavelength spacing is not appropriate. A wavelength plan that spreads the beats around is the only effective way to eliminate 4WM, while securing an effective, robust and high-power transmission in the optical fiber. Such a wavelength plan, as described above, can be enhanced with wavelength offsets, which are especially important for the EML lasers. This is so that they can accommodate the multiple lobes associated with SBS suppression circuitry.

Typically, the recommended maximum total launch power for a well-designed FS wavelength plan is 19 dBm. For 16 wavelengths, light may be launched with up to 7 dBm/wavelength of power launched into the fiber, and for 8 wavelengths light may be launched at 10 dBm/wavelength by the same token. Higher light may be launched for 4 wavelengths, but one must be careful to spread out the wavelengths within accepted guidelines to reduce the effects of SRS and XPM. Since analog wavelengths operate at high power levels, we use the wavelength plans noted above -- and therefore the links between some PHEs and SHEs are congested.

We have seen that the Analog FS wavelengths utilize the wavelength plan described in Figure 8. It is possible that the 'gaps' between the 16 wavelengths can be used by upstream Digital Return Transmitter (DRT) wavelengths. In fact, this is a way to enhance fiber efficiency, because the bi-directional transmission of light utilizes more wavelengths, without the consequent increases in optical intensity at any one end. Thus, we could have 16 WLs in the DS and a like number in the upstream with no significant degradation to the signal integrity in either direction.

The desire to use the same fiber for analog and digital (baseband) operation is an old one -- we first wrote about this in 2006 (SCTE ET). Some of the recommendations there still hold true (although that was about 1 GbE and for CWDM operations). The main idea is to ensure that the analog and digital (in this case 10 Gbps) wavelengths behave as good neighbors and do not influence each other. This is possible when we maintain a substantial optical level differential.

As mentioned, analog transmission is generally at very high levels since the performance needed is quite high. Not so for digital transmission, because the optical input requirement is really in the -21/-22dBm range for signal recovery. This being the case, there is really no need for high power transmission especially if the reach is limited (and we have seen that our links are maxed out at 35km). Therefore, one could launch light at much lower power, which would also reduce the contribution to optical non-linearities. For this reason, if we maintain a substantial optical level differential between the analog and digital transmission, the two formats would behave as good neighbors and not mutually affect each other.




**Figure 12 – Illustrating Analog and Digital Coexistence on the Same Fiber**

Figure 9 shows a representative wavelength plan as seen on the CPM at the SHE, that minimizes non-linearities and enables 4 DS and US wavelengths along with 4 10G circuits on one single fiber. This type of wavelength selection and optical level differential will allow for several digital wavelengths to co-propagate with analog wavelengths. The counter propagation is considerably easier because the return wavelengths are all baseband digital and require very little power compared to the downstream.



## 5. A Word about Coherent Optics

Wavelength planning for accommodating Coherent Optics on the same fiber is an ongoing activity at Comcast and around the industry.



**Figure 13 – Coherent Optics Options**

**Dual Fiber Single Laser Coherent Pluggable:** Traditional Coherent Pluggables are built with an optical engine that has only one laser. The same laser pulls double duty as the transmitting light source that then passes through a nested modulator to create complex light forms to transmit over the fiber. At the far end, the laser receiver adjusts its wavelength precisely to the incoming light, and the heterodyne receiver deciphers the incoming complex signal and send it over to the DSP. The processor then decodes the signals and streams them in either 10Gbps or more likely 25Gbps binary signals to the outside world. The other half of the aforementioned laser's light is sent over to the nested modulator. Its light is then sent over a separate fiber to the first module, where the lasers will heterodyne and decipher the incoming light. This is perhaps the only place where the OBI is intentionally used to enhance the communication system. It is what makes Coherent systems so interesting and special. As mentioned in Section 4, coexisting in a single fiber means that there is not a second fiber available to be to be used, so this option may not work in those cases.

**Single Fiber Single Laser Coherent Pluggable:** For single fiber co-existence, it is proposed to use an optical circulator to reuse the same wavelength. In this case, the light in an optical circulator can move only one way, therefore the transmit light can only move out of the module while the received light might move only inside the module. With this arrangement the wavelength may be reused. The circulator has good but imperfect isolation and is prone to losing isolation at temperature extremes. While this may not be an issue at the PHE or the SHE, it may be an issue at the node, where temperatures can cycle between -40C to +60C. Furthermore, reflections in the system can also impinge on the available dynamic range and reduce the effectiveness of the coherent receiver, thus limiting the fiber reach or speed.

**Single Fiber Dual Laser Coherent Pluggable:** A dual laser design for the coherent pluggable has two separate lasers, one for light transmission and another to receive light. The corresponding transmitting and receiving lasers, in two separate pluggables across the fiber mutually lock on to each other. Because their sources are independent, they deliver true light bidirectionality over different wavelengths. As a result, this arrangement is not prone to any of the reflection issues described earlier. Previously, these types of devices were expensive and had higher heat dissipation. However, the advent of 7 nm DSP technology has propelled efficient drivers and with efficient lasers and cost-effective optical amplifiers these issues have been resolved. These modules are also now very cost effective and have been realized in (C-Formfactor Pluggable 2) CFP2 form factors.

This discussion on coherent optics is intended to prepare the reader for the notion of a “Switch on a Pole” (SOAP) type of an architecture, that would, with the help of a coherent termination device, transform the capacity of access networks dramatically. In this context, we gratefully note the efforts of CableLabs and its P2P link and Coherent Terminating Device (CTD) specification efforts. There are also multiple efforts underway towards 25G NRZ transmission, based on a 10G SFP platform, and also of a new type of point-to-multipoint (P2MP) architecture using subcarriers. All are trying to bring higher speed optics closer towards the access space.

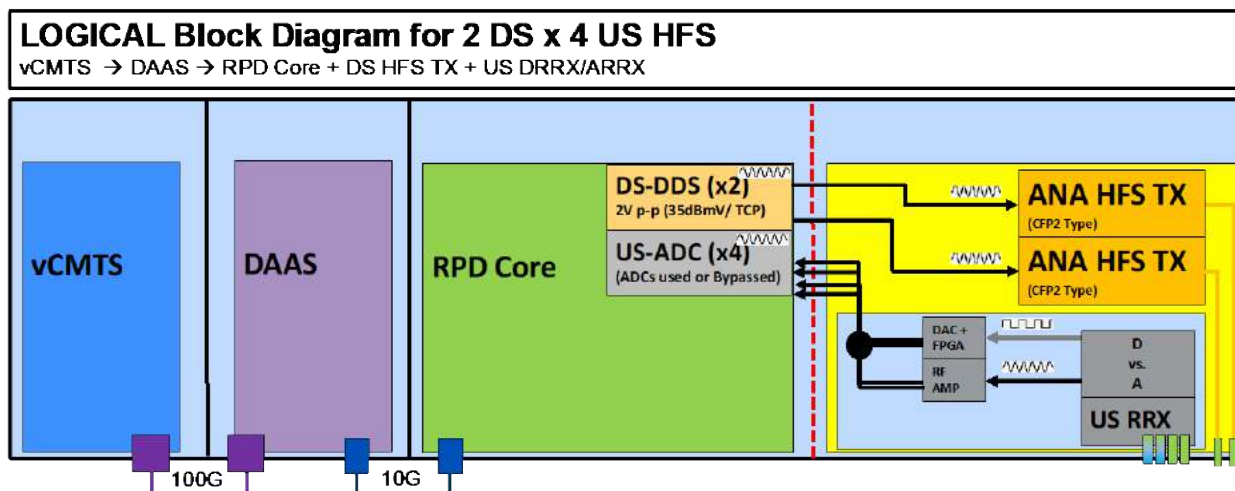
## 6. Hybrid Fiber Shelf

A typical headend that has CMTS is a pretty big facility. Generally, it holds the high-speed circuits coming in to connect to the internet, and video sources coming in either from a satellite dish farm or through terrestrial fiber. Multiple modulators groom the video to be on specific outbound frequencies. Any on-demand servers are connected to modulators as well. The CMTS itself is processing, scheduling and getting traffic, and modulating it. Finally, all of these signals are RF-combined and fed into an analog transmitter.

Each RF combining attenuates the output signal, so a very common requirement is for high RF output power as mentioned in DRFI (DOCSIS RF Interface) Specification to burn through all the of the combining network. Generating high output RF power of sufficient quality requires RF amplifiers and a high-power dissipation to power them up. Of course, all of this is frittered away in the RF combiners until it reaches the transmitter. If there is cabling in a PHE, then the cable loss and its tilt will have to be taken into account. Typical transmitters today require 10 dBmV/6 MHz (in days past they would have required 15 dBmV/6 MHz) but it is common to have individual per channel modulator outputs as high as 60 dBmV/6 MHz.

As we have discussed, there are two main problems with this scenario. For starters, as HSD traffic grows, so too will the CMTS ports or the needed capacity per port. At the limit when HSD has subsumed most of the traffic needs, then the RF can all be potentially available at one port. This will also require the maximum CMTS licenses (discussed in Section 2), but until that time, the entire RF combining infrastructure is carried. However, once the vCOREs come online, and the entire CMTS is virtualized, the timing is perfect to dismantle the RF combiners and recover the space. This is possible because the entire RF spectra is populated by the vCMTS output port and can be connected via a grey 100G connection to the DAAS switch. We note here that the DAAS output port of 10G puts out the complete RF spectra that may be generated at the RPD.

The DAAS may be connected to the headend RPD device thru 10G AOCs, which will have fairly low power consumption. If the RPD is co-located with the VCMTS then this might even be connected via DAC cables that are even more power conserving and at very low cost.



**Figure 14 – Logical Block Diagram for HFS**

A typical RPD without its post amplifier puts out 2V p2p. This is roughly 35dBmV of RF power (for distortion-free performance). Typically, the RPD MER is better than 50 dB, on average.

If a chassis has only RPDs in it, we call it an RFS (for “RPD Shelf”) to distinguish it from a Hybrid Fiber Shelf (HFS), which would have had the transmitters and return receiver integrated. For our purposes, the red line shown in Figure 13 is the demarcation point. We will use the acronym “HFS” generically, to mean that the RPD Shelf and the optics shelf is either integrated or co-located – either way, the RF combining circuits are eliminated or heavily reduced. While the RFS is already of material benefit to us as an operator, the HFS enhances that benefit in cost effectiveness for critical infrastructure.


Years back, the analog transmitter was a device that took considerable space and power and used high performance Directly Modulated Laser (DML) technology. State-of-the-art transmitters are now built using Externally Modulated Laser (EML) technology within the size of a CFP2 form factor typically used for 10G transmission. These comprise an EML, the linearization circuit for it and any SBS suppression circuitry. Typically, we need only around 10-11 dBm of SBS suppression to cover the 30km or so of the SHE to node link. A typical HFS transmitter would have around 4W of dissipation and around 11 dBm of output optical power. As is oftentimes the case, EML transmitters have lower optical modulation depth (OMI) than do the DMLs, but their chirp free dispersion performance and lower RIN make up for it and generally provide a modicum of advantage. The best feature of these is that there is no reason now to declare the fiber lengths.

Using DMLs would have required us to declare fiber distances, because the electronic dispersion compensation would need distance to cancel fiber dispersion-based distortions generated in the fiber. Even though this is a one-time provisioning issue, it causes problems when a redundant route gets activated. As it happens, there is not redundancy from the SHE to the node, but still, eliminating this extra step is a welcome feature. So is the feature of being able to use common optical passives for 10G, coherent and/or analog, which was harder to do with DML transmitters.

In another paper this year, we have described a pervasive monitoring tool to track all wavelengths. That tool also needs a description of the fiber length. One thing to keep in mind here is that the needed number of analog wavelengths per fiber is going to be rather limited in this architecture. In this architecture we typically would have had each SHE feed one node, on average. This is because the link would have


started at the PHE, then thru the SHE and then all the way to the node, and we have seen that long-distance transmission of analog wavelengths is limited.

In fact, while multiwavelength analog transmission is limiting, single or limited wavelength analog transmission over modest fiber distances with EML transmission is has high performance. Figure 15 shows a measured result of an RPD followed by EML transmitter over 38km of fiber. This shows a node MER of ~45dB, which quite handily beats the 4K QAM requirement of 41dB.



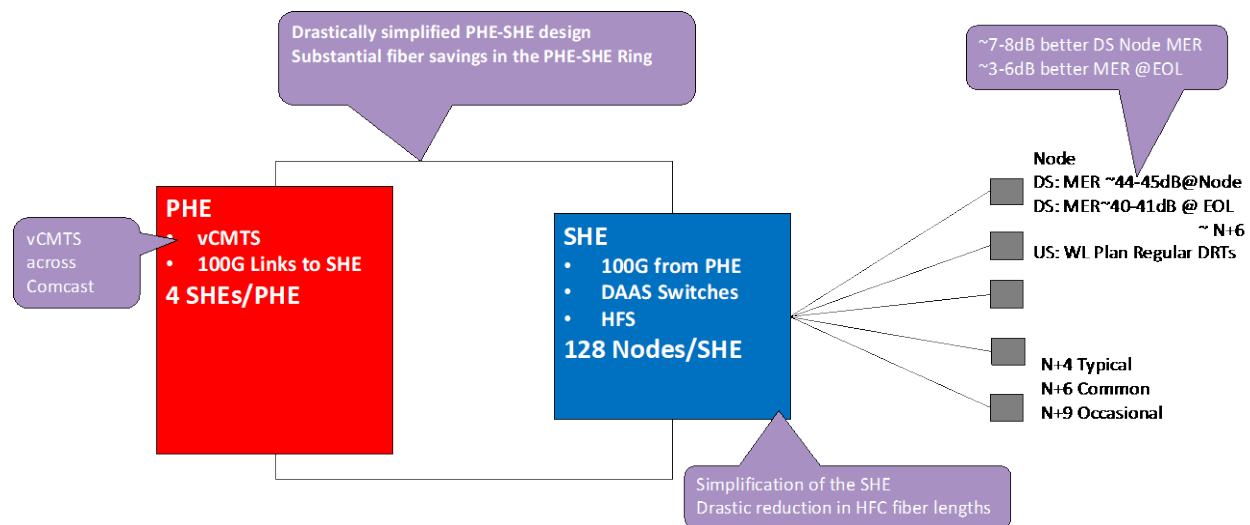
**Figure 15 – Measured RPD and EML Transmitter MER**

However, it is often the case that multiple amplifier cascades are in place after the RPD. In Figure 16, we can see the performance of an amplifier cascade that has been added to the RPD MER. Assuming each RF amplifier at around 50 MER each, this indicates that the cumulative MER performance is also dictated by the RF amplifier cascade as well as that of the RPD and of the HFS EML transmitter. Figure 16 shows a simulated performance of the RF cascade, beginning with an MER appropriately reduced to account for unit to unit variation and temperature performance. The cumulative performance is then very close to ~41dB MER, even for N+6. This performance complements the efforts of DS profile management application (PMA) that automatically adjusts the modulation complexity to available SNR and maximizes capacity deployed at Comcast. With smaller cascades, the performance is even better.



**Figure 16 – Node and RF Cascades**

Figure 17 summarizes the benefits of the Hybrid Fiber Shelf.



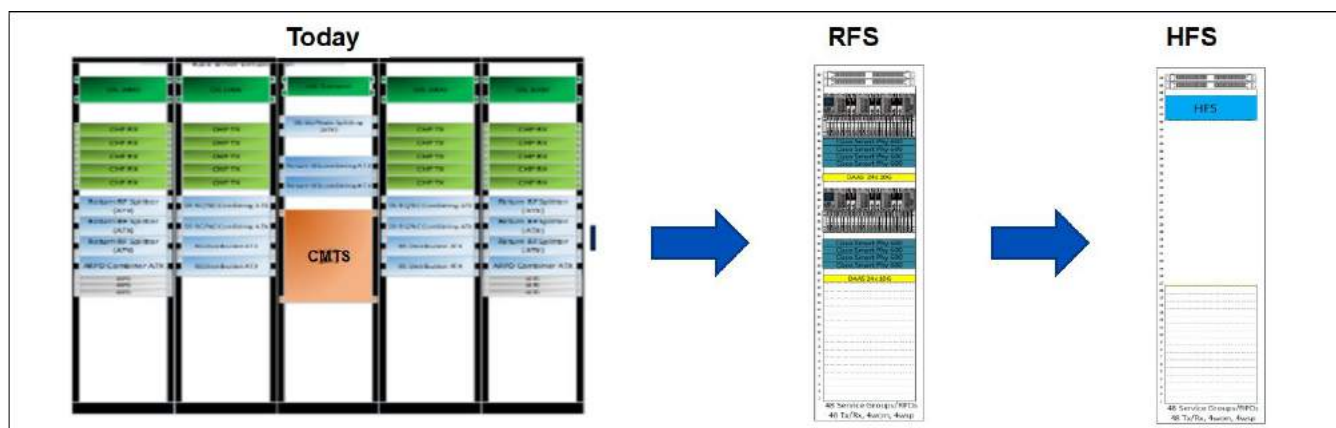
**Figure 17 – Summary of HFS**

## 7. Critical Infrastructure

In this section we talk about how both HFS and RFS can potentially save on Critical Infrastructure (CI) in the PHE and the SHE. As we have indicated before, a regular CMTS occupies several rack units (RUs). Further space is occupied by the video generators, and the most space is occupied by the RF combining circuitry. Overcoming the excess loss of the RF combiners requires that the CMTS and other video equipment put out a rather large RF level. This RF Level is then split and combined to incorporate the various signals and then provided to the Analog Transmitter at an appropriate level. A typical CMTS-based system connected to 192 service groups can take five 7-foot racks worth of valuable headend real

estate. Some of the older CMTSs and the existing RF combining equipment take up even more space, because the CMTS and the DS transmitters are not co-located in the building, so RF cables are run across the headend to connect them. This has the unfortunate result of needing to have slope and amplitude correction circuitry, on a transmitter by transmitter basis, to groom the input RF to the transmitter for distribution. Added to this, because RF combining is complex and folks will not want to touch it once installed, there are numerous test points distributed across the RF cables which provide the convenience of test points but at added power and cost. The same process is repeated for the US as well. The US RF out of the US RX comes in at a high level, is split multiple ways, and combined. This conserves CMTS ports, with enough test ports to track RF levels along, and is then connected up to the CMTS.

With a system described above, accommodating necessary node splits will require additional headend space that is hard to get. Adding to a headend is quite expensive and may not even be possible, in the time needed for the node splits, and after accounting for permits and construction.



**Figure 18 – HFS Critical Infrastructure**

With an enterprise wide vCMTS/vCORE one can now compress all of the HSD, video and other assets into one converged platform and have a single RF connection. In Figure 18, we have indicated that the RFS is co-located. Here, the RF combining is totally replaced with just a DAAS port that is connected to the RFS via AOCs or DACs. Then, there is a modicum of RF circuitry before the output level is connected to the analog chassis. Thus, the space and power savings are considerable, as shown above. Even though the RFS is very compact, it still has a higher output, because we may need to feed existing analog transmitters which may require higher outputs. Sometimes the analog gear may not be co-located (especially in the PHEs) and therefore it may need to be connected to long sections of RF cable. For this reason, the RF output of the RFS is also higher and will require the consequent grooming. One additional fact will need to be considered as well: any time a new node is commissioned, it is common to perform what we call “First Node Provisioning” (FNP). This is done by connecting a CM directly to the chassis and ensuring that the CM can turn up and stay connected. This is done to ensure that there are no loose ends in the node commissioning process. It requires test points for the US that will be able to not only display the ingoing RF levels, but also provide a path for the CM signal to be injected. For this reason, the modicum of RF combining is also maintained.

A study by our Critical Infrastructure group on one of the PHEs found not only substantial space savings, but also that there would be 23% reduction in powering. A 23% reduction in powering is quite a major achievement, because it will also contribute to lower air conditioning costs and our corporate sustainability goals. There is a substantial overall cost reduction with a move to the vCMTS/RFS combining, as well, but that is out of scope for this paper.

The real breakthrough, in hindsight, was the move to a HFS: Integrating the whole Remote PHY device (RPD), downstream transmitter and upstream receiver, all in the same module and installed in the same chassis. When done this way, there is no RF combining, by definition, the test points are conducive not only to verify RF levels, but also to do FNP. Also, the RF level out of the RPD is exactly equal to the RF level needed by the DS transmitter, and the RF level out of the US RX is also exactly equal to the RF level needed for the RPD. In this case, the power requirement is much lower, while the overall density is higher. Note that the RFS will have higher density, but when combined with the need to add in the DS TXs, US RXs and the RF combining, it still results in an overall lower density than the HFS. Furthermore, since the transmitters are very low in power consumption, similar analysis indicated above shows a further 40% reduction in power and space. This makes it very attractive and enables us to start placing the HFS in secondary headends and other such “hard to fit” areas. This is where the maximum benefit of this architecture comes into effect.

It is important to realize that both RFS and HFS have their place. For example, in a PHE, where there is a preponderance of existing optical transmitters and there is a modicum of space, many of these could be repurposed and either co-located or distance located from the RFS. But if the issue is to move closer to the nodes or if node-splits are to happen and new transmitters and receivers are needed anyway, then a move to SHE may be warranted and here the HFS would be a far better fit.

It may be asked that if additional nodes are to be added, why would we not just deploy RPD nodes. Indeed, that may be a legitimate option, and 1x2 and 2x4 RPD nodes are more available now than they were a couple of years back. However, if the aim was predominantly to eliminate traditional CMTS and standardize on the vCMTS, replacing of all nodes would be too onerous of a task. A HFS would be a better choice and enable use of existing nodes with available segmentation capacity. We have further shown that the performance is quite good and holds itself well in a N+x type system, where there is a fair contribution from the RF cascade as well.

To summarize, a combination of DAA RPD nodes, 1x2, 2x4 HFC RPD nodes, RFS and HFS are a set of tools to accelerate the virtualization of CMTS and usher in uniformity and standardization within our access networks.

## 8. Conclusions

We began this paper with a recap of our work to virtualize what are some very complex portions of our network, including the CMTS, which resulted in network simplification and harmonization of multiple purpose-built platforms into one common entity. However, DAA deployments require headend and field modifications as part of the deployment. RFS takes great step forward by accelerating virtualization with a design of RPDs in a shelf, but still requires some RF combining and a separate optical chassis.

HFS takes that concept further, integrating with virtualized CMTS/DAA switches at one end and with transmit/receive analog optical signals at the other end, into the outside plant. This innovation provides rapid and economically sustainable increases in capacity by accelerating integration with the vCMTS, as DAA construction proceeds in other areas of the network. HFS improves critical infrastructure in headends by saving wasted space, power and time inherent in RF combining and splitting circuits. Extending HFS into secondary headends also enables significant fiber reclamation. Locating these assets in secondary facilities, closer to the traditional HFC nodes, improves performance in ways that translate into enhanced capacity.

## 9. Acknowledgements

It is with gratitude that we acknowledge the team within Comcast who has been working on the RPD and RFS projects. We also thank the Critical Infrastructure team for their analyses, and the Access Engineering team for their support on fiber convergence and planning efforts.

## Abbreviations

|      |                                        |
|------|----------------------------------------|
| 4WM  | Four wave mixing                       |
| AOC  | Active Optical Cable                   |
| CAGR | compound annual growth rate            |
| CFP2 | C-Formfactor Pluggable -2              |
| CI   | Critical infrastructure                |
| CM   | cable modem                            |
| CMTS | cable modem termination system         |
| CPM  | Continuous Pervasive Monitor           |
| CTD  | Coherent Terminating Device            |
| DAA  | Distributed access architecture        |
| DAC  | Digital Access Cable                   |
| DAAS | Distributed access architecture switch |
| DRFI | DOCSIS RF Interface Specification      |
| DRT  | Downstream Return Transmitter          |
| DML  | Directly modulated laser               |
| DS   | Downstream                             |
| DSP  | Digital Signal Processor               |
| EML  | Externally modulated laser             |
| FNP  | First node provisioning                |
| FS   | Full spectrum                          |
| GOA  | Grey Optics Architecture               |
| HFC  | Hybrid fiber coax                      |
| HFS  | Hybrid Fiber Shelf                     |
| HSD  | High speed data                        |
| MAC  | Media access control                   |
| MER  | Modulation error ratio                 |
| MPEG | Moving Pictures Experts Group          |
| MSO  | Multiple systems operator              |
| NRZ  | Non Return to Zero                     |
| OBI  | Optical beat interference              |
| P2P  | Peer to peer                           |
| P2MP | Peer to multi-peer                     |
| PHE  | Primary Headend                        |
| PHY  | Physical                               |
| QAM  | Quadrature Amplitude Modulation        |
| RF   | Radio Frequency                        |
| RFoG | Radio Frequency over Glass             |



|       |                                        |
|-------|----------------------------------------|
| RFS   | RPD Fiber Shelf                        |
| RIN   | Relative intensity noise               |
| RPD   | Remote PHY device                      |
| RU    | Rack unit                              |
| Rx    | Receiver                               |
| SFP   | Small form-factor pluggable            |
| SHE   | Secondary Headend                      |
| SOAP  | Switch on a pole                       |
| SRS   | Stimulated Raman Scattering            |
| US    | Upstream                               |
| vCMTS | Virtual cable modem termination system |
| VOD   | Video on demand                        |
| WL    | Wavelength                             |
| XMP   | Cross phase modulation                 |

## Bibliography & References

1. *Operationalizing the Grey Optics Architecture: An Update a Year After*, Venk Mutalik, Dan Rice, Bob Gaydos, Doug Combs and Pat Wike, SCTE EXPO 2020
2. *It is 10PM: Do you know Where Your Wavelengths are? Continuous and Pervasive Monitoring of Optical Assets in the Access Domain*, Venk Mutalik, Dan Rice, Rick Spanbauer, Simone Capuano, Rob Gonsalves and Bob Gaydos, SCTE EXPO 2020
3. *Distributed Access Architecture – Goals and Methods of Virtualizing Cable Access*, Nagesh Nandiraju et. al., SCTE EXPO 2016
4. *When Wavelengths Collide, Chaos Ensues: Engineering Stable and Robust Full Spectrum Multi-wavelength HFC Networks*, Venk Mutalik et. al., SCTE Cable-TEC EXPO 2011
5. *Gigabit Ethernet and Analog/QAM Traffic Compatibility in HFC Networks: A case for Physical Layer Convergence*, Shamim Akhtar, Doug Weiss, Venk Mutalik, Marcel Schemmann, David Heisler, Down Wesson, Liyan Zhang, SCTE ET 2006

# Operationalizing the Grey Optics Architecture:

## An Update One Year After

A Technical Paper prepared for SCTE•ISBE by

**Venk Mutalik**

Executive Director, HFC Architecture  
Comcast  
1401 Wynkoop St, Denver, CO  
+1 (860)-262-4479  
Venk\_Mutalik@Comcast.com

**Dan Rice, Comcast**

VP, HFC Architecture  
1401 Wynkoop St, Denver, CO  
Daniel\_Rice4@Comcast.com

**Bob Gaydos, Comcast**

Comcast Fellow  
1800 Arch Street, Philadelphia, PA  
Robert\_Gaydos@Comcast.com

**Doug Combs, Comcast**

Consultant Engineer  
1401 Wynkoop St, Denver, CO  
Doug.Combs@Comcast.com

**Pat Wike, Comcast**

Sr. Director, Access Engineering  
1401 Wynkoop St, Denver, CO  
Patrick\_Wike@Comcast.com

# Table of Contents

| Title                                                          | Page Number |
|----------------------------------------------------------------|-------------|
| 1. Introduction.....                                           | 3           |
| 2. Architecture Recap.....                                     | 3           |
| 3. Initial Deployment Site .....                               | 6           |
| 4. Initial Deployment.....                                     | 8           |
| 5. Heart to Limbs .....                                        | 10          |
| 6. Deployment Guidelines .....                                 | 12          |
| 7. Ingress Identification and Mitigation.....                  | 13          |
| 8. Tooling, Eventing, Layering ... All in a Sprint.....        | 15          |
| 9. Provisioning the GOA-GOT system .....                       | 17          |
| 10. “Grafana-ing” the Constellation .....                      | 18          |
| 11. Orchestration and Association .....                        | 20          |
| 12. Secure, Remote and Automatic: RPD Life Cycle Manager ..... | 21          |
| 13. Eventing and Ticketing.....                                | 22          |
| 14. Future Steps.....                                          | 23          |
| 15. Acknowledgements .....                                     | 23          |
| Abbreviations .....                                            | 24          |
| Bibliography & References.....                                 | 24          |

## List of Figures

| Title                                                                                      | Page Number |
|--------------------------------------------------------------------------------------------|-------------|
| Figure 1 – Traditional DAA in Primary and Secondary Headends and Plant .....               | 4           |
| Figure 2 – Illustrating GOA Architecture .....                                             | 4           |
| Figure 3 – Illustrating all Nodes and their Power Dissipations.....                        | 5           |
| Figure 4 – Upgrade Models: Localized Growth (Left) and Uniform Growth (Right).....         | 5           |
| Figure 5 – Technical Details of Initial Deployment Area .....                              | 7           |
| Figure 6 – Geographical area illustrating the Physical Subdivisions .....                  | 8           |
| Figure 7 – GOA and GOTs assembled in the headend, Closed (Left) - Open (right).....        | 9           |
| Figure 8 – Splice Matrix and the Wiring Diagram .....                                      | 11          |
| Figure 9 – Illustrating the Basic CLI and how it summarizes the GOA Constellation.....     | 13          |
| Figure 10 – CLI with Various Settings.....                                                 | 14          |
| Figure 11 – Highly Simplified Ingress Illustration and the Role of US RF Attenuation ..... | 15          |
| Figure 12 – The Tooling Diagram – Cloud Based Provisioning/Monitoring/Tooling.....         | 16          |
| Figure 13 – Visualizing All US and DS Metric at a Glance Across the Node .....             | 17          |
| Figure 14 – Provisioning Application Flow .....                                            | 18          |
| Figure 15 – Grafana and Real Time and Historical Constellation Information .....           | 19          |
| Figure 16 – Illustrating the concept of Association and Orchestration .....                | 20          |
| Figure 17 – Association and Orchestration in the GOA RPD .....                             | 22          |
| Figure 18 – Automatic Eventing and Ticketing .....                                         | 23          |

## 1. Introduction

High Speed Data (HSD) customer growth and capacity requirements continue to increase year over year. A Distributed Access Architecture (DAA) provides enhanced upstream and downstream capacity to match demand. Last year, we covered a new approach in DAA called Grey Optics Aggregation (GOA) that provided significant savings and enables DAA deployment at scale regardless of plant density.

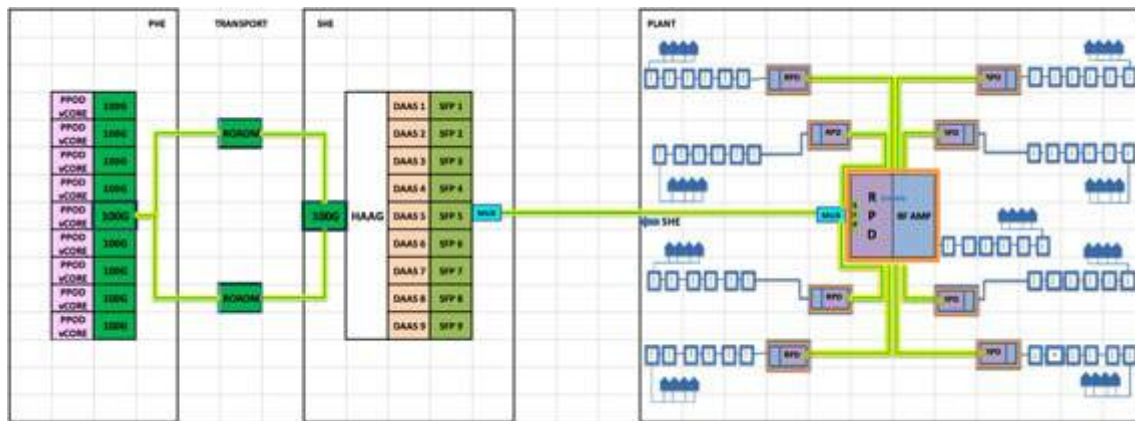
This year, we report on the lessons learned during our GOA trials and deployments of the architecture. In addition to proving out the technology and cost benefits in the field and observing its stability, deploying the architecture also enabled us to interconnect and create an ecosystem of remote monitoring and control for the various GOA elements. We will report on our ability to continuously monitor Grey Optical Terminating (GOT) nodes, detect cable modems connected to specific GOT nodes and remotely identify and mitigate ingress in the GOA domain.

Many of the remote techniques were used during the COVID pandemic and can potentially be generalized further to enable remote monitoring and mitigating capabilities across the network. Furthermore, we report on recent work on bi-directional Coherent Optical modules that enable efficient extension of the architecture, which, over a period of time, can create cost effective convergent and scalable “Switch On A Pole” (SOAP) access networks.

## 2. Architecture Recap

Last year, we presented a new Grey Optics Aggregation (GOA) architecture that can reduce cost in the outside plant, reduce cost and space in primary and secondary headends, and provide substantial upgrade opportunities to increase capacity and enable new Switch On A Pole (SOAP) type architectures. The GOA architecture works especially well in medium and lower density areas and can exist side by side with traditional nodes equipped with Remote PHY Devices (RPDs). As such it helps to accelerate the virtualization of the CMTS (vCMTS) by enabling it to be deployed across complete primary and secondary headends.

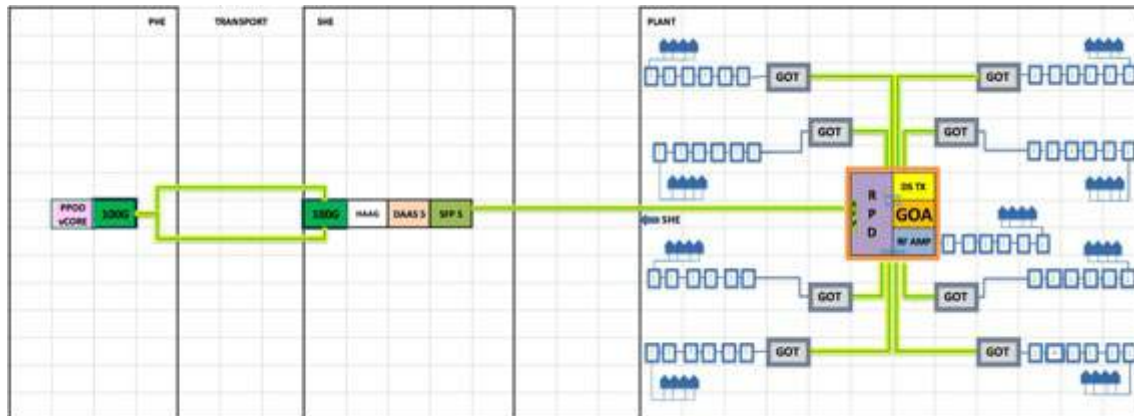
Figure 1 shows a simplified DAA deployment of vCMTS, DAA Switch (also called DAAS, connects vCMTS data to RPDs) ports and RPDs in the field. Recall that each RPD requires its own DWDM SFP (Dense Wave Division Multiplexer Small Form Factor Pluggable transceiver), multiple such SFPs can be optically multiplexed and connected to individual DWDM SFPs in DAAS ports through optical de-multiplexers. Multiple DAAS ports are aggregated and connected to the vCMTS core via optical methods common to long haul networks (~100G -400G).



**Figure 1 – Traditional DAA in Primary and Secondary Headends and Plant**

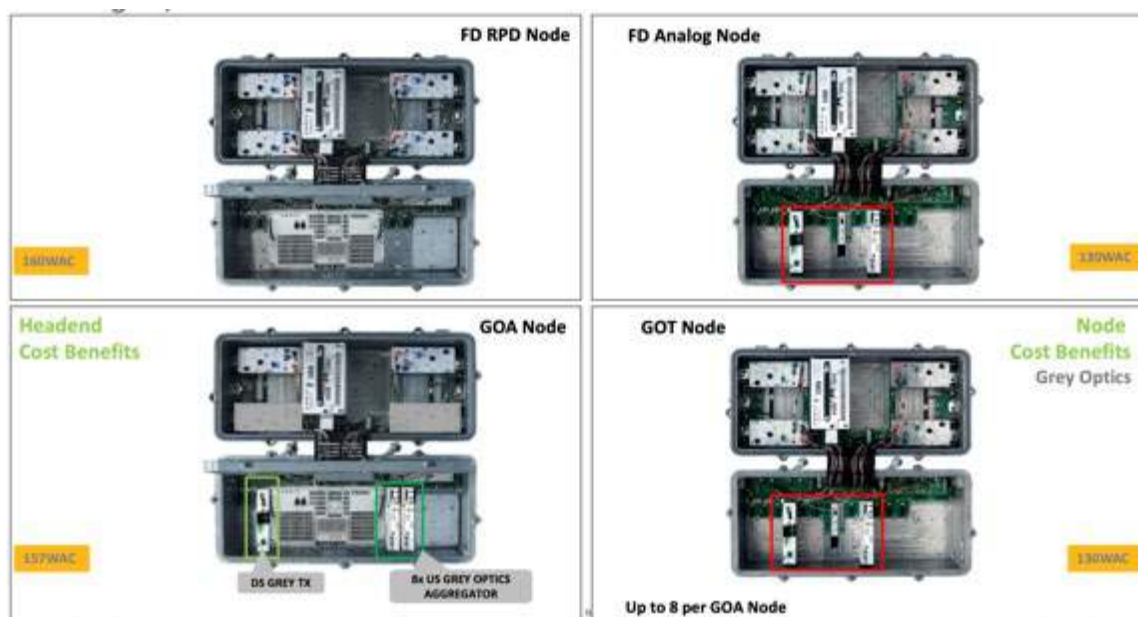
This architecture works well when home/neighborhood densities are high and as such the cost of each RPD and the associated optics at the headends can be spread across a large base of household passings (HHPs). However, if the area is sparsely populated, this may be difficult. Even when a high-density residential area is selected, it can contain pockets of lower density. Addressing them cost-wise might be the proverbial tail wagging the dog.

Presented in Figure 2 is the GOA Architecture that provides a cost effective and pervasive way to deploy virtualized CMTSs across the whole DAA footprint.



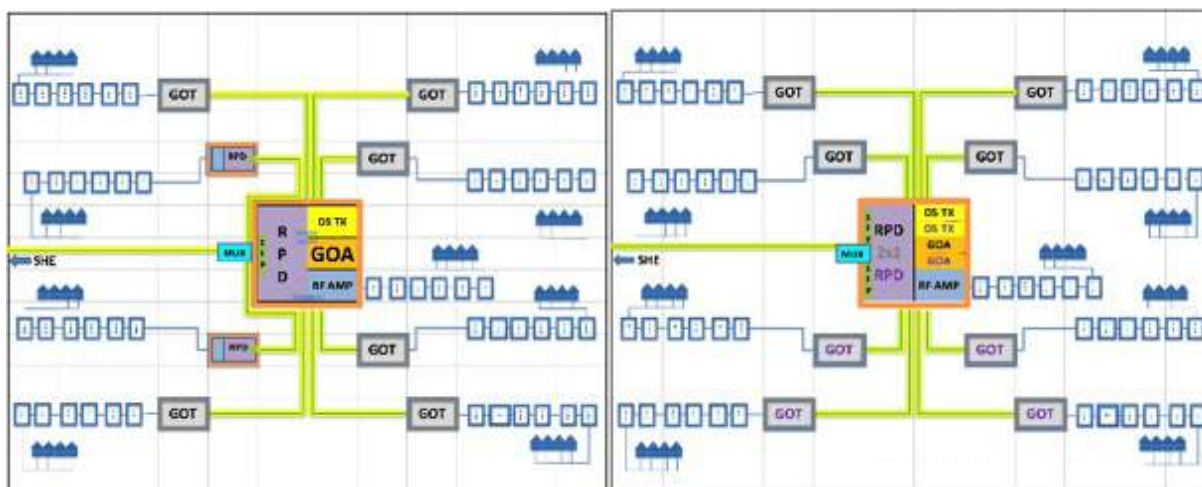
**Figure 2 – Illustrating GOA Architecture**

In this architecture, a prime node contains the RPD and is called the GOA (Grey Optics Aggregating) node, since it is populated with an analog forward transmitter and return receiver modules for aggregation. The downstream (DS) transmitter output is then split and connected to multiple (up to 8) nodes, called Grey Optics Terminating (GOT) nodes. Each of the GOT nodes contains a downstream receiver module and distributes the RF output to its four constituent RF blocks. The RF levels of the GOT node both for downstream and upstream (US) traffic are identical to regular RPD nodes. For the US, the RF signals get combined in the GOT RF tray and fed into a digital return transmitter (DRT). From there, the signals are sent to the GOA node, where they are fed to the aggregating receiver module, one each for each receiver. Then they are RF-combined and fed to the RPD, thus completing the signal circuit.



**Figure 3 – Illustrating all Nodes and their Power Dissipations**

As can be seen in Figure 3, all nodes are the same physical size. And it is the case that all of the above nodes have identical RF US and DS characteristics. The GOT node has an optical AGC receiver that always adjusts the RF level of the receiver to be identical to the RPD output. The DRT to the aggregate return receiver module gain is set to unity, so that the combination of the DS receiver and the US DRT are an almost perfect replacement for the RPD RF levels. In fact, it is for this reason that replacing the DS receiver, the US DRT and the node monitoring card together with an RPD would instantly transfer the GOT into a full RPD node when and if needed.



**Figure 4 – Upgrade Models: Localized Growth (Left) and Uniform Growth (Right)**

Figure 4 illustrates one of the major advantages of the GOA architecture: Not only does the GOA allow for cost effective fiber deployments across multiple areas, but it also allows for easy upgrades, should the need arise. Generally, the DAA deployments are projected to be stable and address growth needs for a substantial period of time. However, if there is localized growth at any one of the GOTs that defied earlier

predictions, the move from a GOT node to the GOA node is easy, from a field perspective, because it just replaces the analog receiver and DRT with the RPD and connects the two ends to the optical passive. Furthermore, the new RPDs that would need to be deployed would be more computationally capable, consume lower power and occupy less space (as described by Moore's Law, Dennard's Law and Koomey's Law, respectively) and cost less as well. This will, in all likelihood, enable us to use the uniform growth model (shown in Figure 4) to deploy a 2x2 RPD that provides increased capacity across the board, when the time comes thus requiring far fewer site visits for network upgrades. Furthermore, a move to Full Duplex DOCSIS (FDX) or Switch on a Pole (SOAP) that relies quite heavily on RPD, Node and vCMTS enhancements would also be well served by the GOA architecture.

For all of the above reasons, the GOA architecture was deployed on a trial basis and its performance, robustness, tech friendliness and customer satisfaction have been established. With this information, deployments are now occurring in additional areas.

In this paper, we discuss some of the lessons learned during the deployment. In addition to the above features, several innovative features suggested themselves. One was an intelligent use of the aggregation modules to associate cable modems (CMs) not only to the RPD and to the vCMTS core, but also to each of the GOTs. In addition, the same type of intelligence enables us to identify ingress or noise in each of the GOT nodes or on the RF blocks in the GOA node. These features and other interesting applications of harnessing the monitoring information flow between the GOTs and the GOA will also be discussed.

### **3. Initial Deployment Site**

After discussions, a suburban area in one of the Comcast divisions was chosen for initial deployment. This area already had just been built with traditional RPD nodes. Figure 5 is a description of the chosen area, which served ~550 HHPs with 21 RPD nodes. Some nodes served as few as 15 HHPs while the maximum HHP count was 40, with an average count of 25 HHP / RPD.





**Figure 5 – Technical Details of Initial Deployment Area**

As shown in Figure 6, the selected area was a regular subdivision, as would be expected off of a major metropolitan area. What we found was that the common RF build practice of running RF through the front of each property meant that the number of HHPs that could be served per node was rather limited. This is especially true given that DAA builds follow a node plus zero amplifier (N+0) design, thus do not have amplifiers to extend RF runs.





**Figure 6 – Geographical area illustrating the Physical Subdivisions**

Because of this, it is easy to see that a build that spans across a SHE location may have on average have high HHP/DAA Node density but could still have several nodes with much lower HHP/DAA node densities. We have investigated many ways to enhance the HHP/Node, including a thorough analysis of the drop levels needed. A judicious application of those design rules does extend the number of HHPs/Node and are defined as acceptable in newer playbooks.

In this case, taking a look at the geography enabled us to see that the 21 RPD nodes could be reduced to 3 GOAs and 18 GOT nodes. This represents a 7x reduction in amount of SFPs, DAAS ports and vCMTS microservice instances required. The GOA locations are indicated in Figure 5, along with the GOT nodes and their HHP counts.

## **4. Initial Deployment**

To prep for the transition from RPD to GOA, we set up a team of analytics and traffic experts. We then charted all the CMs associated with each of the nodes and noted the cumulative traffic in each RPD. Then the team threaded the total traffic and time series, aligning them to estimate the cumulative traffic, were the nodes converted to the GOA architecture.

A word about HHPs and CMs. While there is wide variation in respect to broadband penetration rates and number of DOCSIS devices per HHP, in general, when both are taken together, the number of CMs in a given area is approximately equal to the HHPs in that same area (this was certainly true for our case here). With this in mind, we will use HHP and CM count interchangeably in this paper.

We did this to get a baseline view of the traffic before the RPD nodes would be aggregated with the GOA architecture. This laborious process was then automated, and we found that as the aggregated node size increases, the cumulative traffic increases, but at a slower rate than the increases in HHP. And, as expected, we find that the traffic, which is lumpy with small numbers of CMs, gets more blended with larger numbers of CMs. The idea was to ensure that none of the GOAs would be overwhelmed with traffic -- and none were.

We also put together a dashboard that threaded the US RF output levels and input levels, SNR, PER/FEC rates and RF frequency per CM, as well as DS input levels and MER for each frequency, from multiple sources. Some data were directly from the combined vCMTS/RPD and others from the CMs, as they reported back. All of these were time stamped and reported daily as Max, Min, Mean and 95<sup>th</sup> percentiles. This enabled us to be able to see each CM and all associated metrics simultaneously to see the impacts in almost real time. Over time, this was very useful.

Once the decision to move ahead was taken, the division leaders took over the process of combining nodes in an appropriate manner and drawing up a schedule for node cutovers. The leadership, engineering talent and commitment of the division was unquestionably paramount in the move forward. As can be seen from the figures above,

- GOA A had 240 HHP, with 8 GOT nodes
- GOA B had 202 HHP, with 7 GOT nodes
- GOA C had 109 HHP, with 3 GOT nodes

These were selected in this fashion to ensure that we had sufficient opportunities to verify performance with larger and smaller HHP groups and GOT node aggregation. In parallel, we also performed lab tests to determine that the RPDs could handle the 240 HHPs that were likely to be added. In reality, the RPD and vCMTS combination can handle much more in some of the newer virtualized systems.



**Figure 7 – GOA and GOTs assembled in the headend, Closed (Left) - Open (right)**

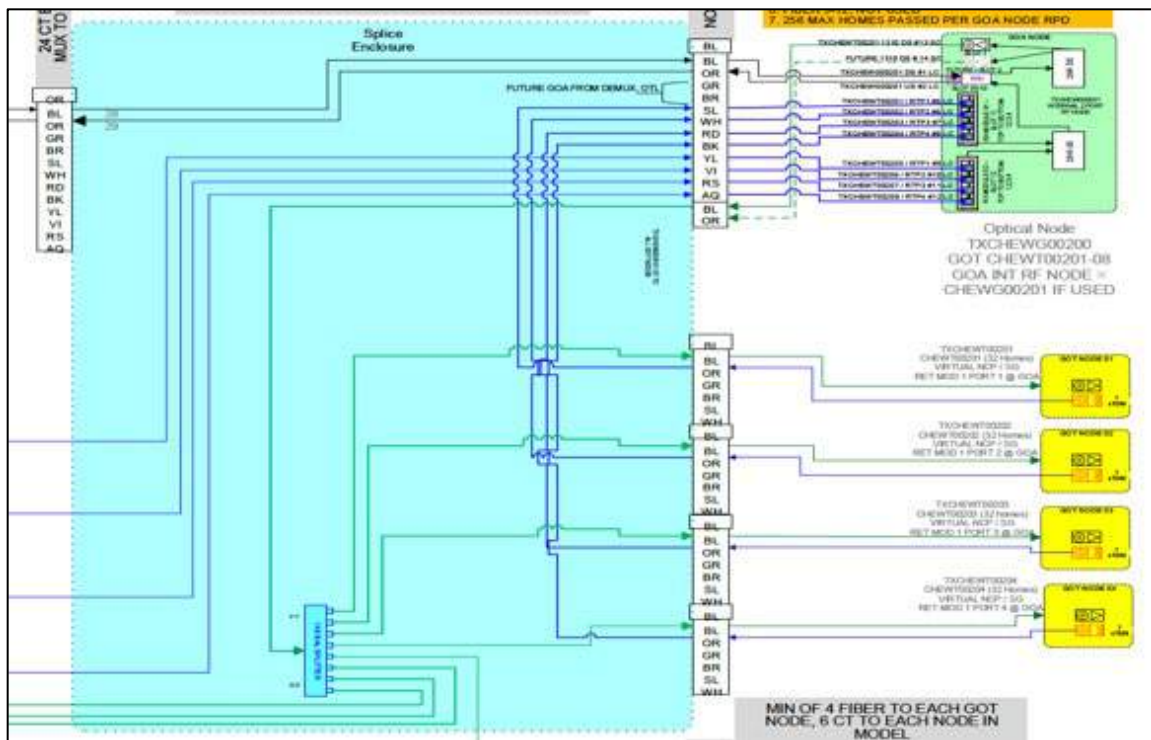
Once this was decided, we drew up entry and exit criteria and a backout plan. Part of the plan was to train our field technicians on the new products, and also to set up a full GOA-GOT system in the headend, to ensure that we had a working system for burn-in, and to work out any kinks. Setting up the GOA-GOT system allowed us to verify some of the very preliminary software features then available. These were road-mapped to improve over time and were deployed and tested on the headend system.

When it came time for construction, in a maintenance window, the first GOA was deployed. After a day of observing its performance, the GOTs were then cut in during daytime. This is because the RPDs have a turn-on time, while the GOT nodes (being analog) are turned on instantly with no additional settle-in time needed. All in all, the 21 RPD nodes were turned in to 3 GOAs and 18 GOTs with careful planning, diligent effort led by the division, and periodic checks on customer call logs and performance verifications, using available internal tools, which will be described in detail later in Section 7.

Overall, the construction of the GOAs and GOT (i.e migration from RPDs to GOA) was uneventful, thanks in large part to the planning and efforts described above. The performance of the units was well within what was predicted, there were no customer impacting events, and the customer experience was maintained well. Next we will characterize the many additional software and hardware innovations that were developed and deployed to enhance the useability of the platform.

## **5. Heart to Limbs**

Construction of the GOA architecture follows a “Heart to Limbs” strategy. By that we mean that the GOA node is first deployed. Its functionality is established, the local splice enclosure is primed and then the GOT nodes are activated, one after another. Since the RPD in the GOA node feeds all other nodes, this strategy will require that the plant that feeds all other GOT nodes stays online even after the GOA is deployed. This is not the case for RPD deployments, since each RPD node is individually feeding its own limited set of HHPs. To achieve good visibility and to speed up the construction process, the team came up with an innovative splicing matrix approach, depicted in Figure 8.



**Figure 8 – Splice Matrix and the Wiring Diagram**

In the case of the trial deployment, the fiber plant was already designed with RPDs in mind. In cases where GOA is the preferred deployment, we allocate a minimum of 4 fibers per GOT (typically 6 fibers). The GOA is allocated a 14-fiber pigtail. The standard 12 colors, from Black to Aqua and Black and Orange, repeated. For any GOT, the two fibers that are spliced are color coded so that the Black is always connected to the DS receiver and the Orange is always connected to the DRT. On the GOA side, the GOT designated GOTs 1 – 8 follow the Slate – Aqua colors of the standard Fiber Color scheme. For the GOA RPD, the Black is connected to the Rx of the SFP, and the Orange is connected to the Tx of the SFP. The other Black is connected to the DS Tx and goes to the 8-way DS splitter and then on to the GOT's DS Receiver.

We also followed a minimal design for optical passives that connect to the GOA RPD. We generally note that a parent node rarely has more than 3 GOAs. Therefore, the DAA Mux has been standardized to the 12 port Mux (ITU 61 – ITU 50) and exclusively uses sub-band tunable SFPs for the DAAS ports and for the GOA RPDs.

The GOA architecture, in addition to performing well as a cost reduced DAA option, also contributes to the elevation of DAA architectures as a whole. The GOA node has an I2C bus (this is a bus that enables protocols that can monitor or control various elements inside the node) that connects the RPD to the aggregating Receiver modules, and reports the state of DS receiver power, the US transmit power, the US receive power and the DS transmit power (more on this in Section 10.) But it also connects the RPD to the node itself and helps report on the RF amplifier blocks and the overall power and current usage of the node. Furthermore, each of the US aggregating receivers can either attenuate their own RF out by 6dB or turn themselves off. This helps identify ingress sources amongst each of the GOT nodes. This same feature is available in each of the RF blocks, as well, where the RF block can attenuate the US by 6dB or

turn it off. This feature is called the Ingress Control Switch (ICS). As an aside, we currently do not use these features in regular RPD nodes. The primary reason was the lack of available remote software control or monitoring. With the software development that took place for the GOA architecture use of I2C/ICS can now be used in regular RPD nodes as well in monitoring and ingress detection on the node ports. With this understanding, most of the software that is described in the paper can apply equally well to GOA architecture or to normal RPD deployments.

## 6. Deployment Guidelines

In terms of plant makeup, roughly 60% of our plant is aerial and 40% underground. While DAA is a preferred architecture for high growth tier 1 markets, and it helps us stay ahead of impending node splits and also prepare the plant for further DS and especially US capacity, it takes a fair amount of work to lay the fiber cable, optimize the RF cabling, condition the taps and light up the nodes. In recent times, we have made good progress on optimizing the various parameters that go into designing the network, to improve its efficiency and to select areas especially well fitted for DAA. As reality goes, this is easier said than done – which in turn makes the GOA architecture a good fit for enabling the DAA to be more successful in moving to the vCMTS architecture. To this end, a set of rules have been established for the GOA deployments:

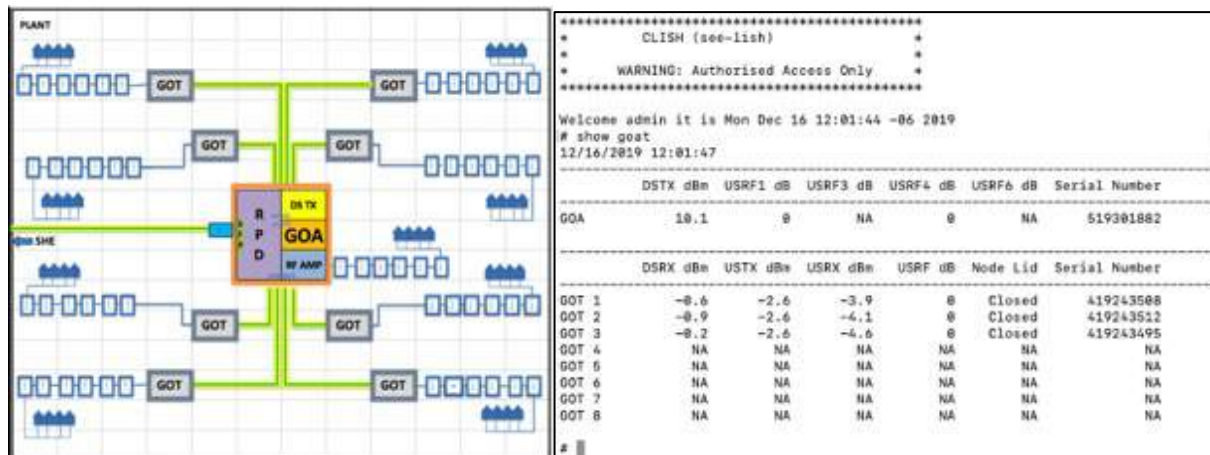
- The GOA node has 2 RF ports and operates at standard DAA RF levels (73.6 dBmV TCP for 1.2 GHz load for the DS and 8 dBmV for the 85MHz US)
- The GOT node has 4 RF Blocks and operates at standard DAA RF levels (73.6 dBmV TCP for 1.2 GHz load for the DS and 8 dBmV for the 85 MHz US)
- The GOA the GOT and the RPD nodes use the I2C and the ICS functionality to communicate/monitor and identify/mitigate ingress
- The GOA RPD has a maximum of 256 HHP, this is the sum total of all passings of the GOA and all associated GOT nodes (in rare instances, an additional 30 HHP may be added to the total count if it can avoid having to add a GOA altogether)
- No more than 8 GOT nodes may be subtended by a GOA node, and it is highly preferred that the GOT nodes be contiguous (in other words, a GOA and its constellations should be in proximity rather than be interspersed with GOTs of other GOA nodes)
- A power interruption extension circuit must be deployed with the GOA RPD (this is a recommended for all RPDs, not just a GOA RPD), and no such power pack is required for the GOT nodes (since they are instant-on nodes)
- 160W is designated per node for dissipation, even though GOT nodes dissipate considerably less. This is to be prepared for when GOT nodes need to be upgraded to an RPD node.
- As discussed, the optical passives deployment uses 12 wavelength devices (ITU 61-ITU 50) and use the sub-band tunable SFPs for GOAs and the DAAS ports. The optical splitter will always be an 1x8 device
- Regular DAA nodes and GOA nodes can be used in the same build footprint at convenient locations, as needed, and the DAAS ports and the vCMTS continue functioning with these devices

These guidelines were presented to our design partners, and with their help we have elected to deploy the GOA architecture in two of our divisions this year (this is in addition to the GOA deployment discussed in section 2), covering multiple DAA builds. “Fingers are crossed” to achieve this, given the COVID pandemic and its impact on construction schedules. It is recognized that underground construction is especially challenging in the current environment and further thought is being directed towards this end.



## 7. Ingress Identification and Mitigation

As discussed earlier, all GOA and GOT nodes have the I2C connectivity to monitor and control many node components. The RPD itself is controlled by Command Line Interface (CLI) commands. There is a secure way to access the RPD to be able to determine how the GOA is configured and display it.



**Figure 9 – Illustrating the Basic CLI and how it summarizes the GOA Constellation**

Access to the CLI for a GOA RPD displays the entire constellation of GOT nodes connected to the GOA. All the GOT nodes that are connected to the GOA have a one-way EMS connection to the GOA via the DRT. Therefore, the GOA receiver is aware at all times about the GOTs connected to it. The GOT node sends its state periodically to the GOA, which then updates this information and passes this on to the vCMTS. In addition, this information is also available at the RPD and can be uncovered via the CLI described above. Section 10 will describe how the information that arrives at the vCMTS can be used, but for now we describe two important aspects of what the GOA can provide which could be of great value.

The entire GOA constellation is represented in Figure 9. The output power of the DS transmitter (DSTX) is displayed and it is nominally 10 dBm. Next, for each of the GOT nodes, the DS receiver power (DSRX) is displayed, which is nominally 0 dBm to -1 dBm, because the DS light passes through an 8-way splitter and traverses some fiber. The US Transmitter (USTX) power is shown next, which is nominally -3dBm (this is a lower-cost Grey 1310 nm SFP). The US Receiver (USRX) power reported by the GOA is shown next, and is nominally -4 dBm, because the US light does not go through any splitter or combiner.

In addition to the light levels, we mentioned earlier that there is an ICS; each of the states of each ICS are reported as well. Thus, the USRF ports refer to the attenuation setting of the ICS of each of the RF blocks in the GOA (recall that we have 2 RF blocks per GOA). Recall that the GOA US Digital Receiver has the capability to attenuate the RF output of each individual GOT US as well. That state is also represented in Figure 9. Furthermore, each GOA and GOT serial number is reported as well. As a neat feature, the GOT node lid is reported as open or closed.

With these settings established, we can now remotely monitor and control the whole constellation. We will show in Section 8 how we have automated this, and how a ticketing structure was developed to automatically indicate major and critical alarms and identify escalation and mitigation options.

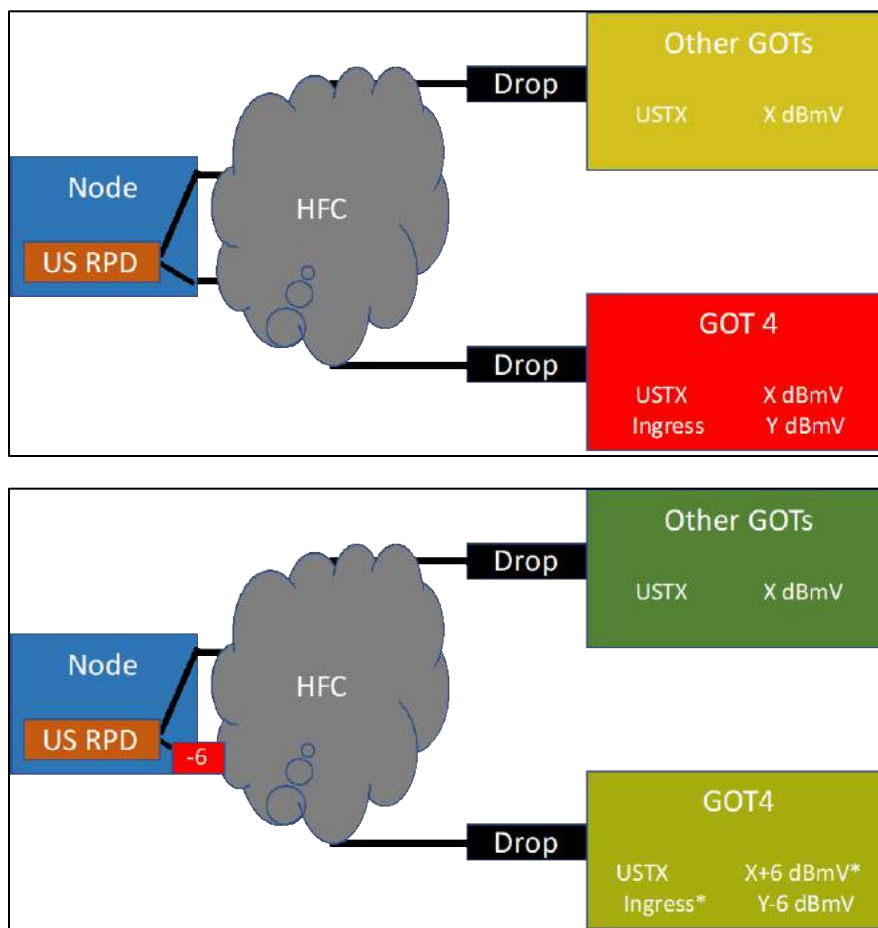
|                     |          |          |          |          |          |               |
|---------------------|----------|----------|----------|----------|----------|---------------|
| # show goat         |          |          |          |          |          |               |
| 11/20/2019 16:24:47 |          |          |          |          |          |               |
|                     | DSTX dBm | USRF1 dB | USRF3 dB | USRF4 dB | USRF6 dB | Serial Number |
| GOA                 | 10.0     | 6        | NA       | 0        | NA       | 519301888     |
|                     | DSRX dBm | USTX dBm | USRX dBm | USRF dB  | Node Lid | Serial Number |
| GOT 1               | NA       | NA       | NA       | NA       | NA       | NA            |
| GOT 2               | 0.2      | -2.7     | -9.0     | 6        | Closed   | 419243545     |
| GOT 3               | NA       | NA       | NA       | NA       | NA       | NA            |
| GOT 4               | -0.2     | -2.6     | -9.1     | 0        | Closed   | 419243566     |
| GOT 5               | NA       | NA       | NA       | NA       | NA       | NA            |
| GOT 6               | 0.2      | -2.6     | -8.9     | 0        | Closed   | 419243513     |
| GOT 7               | NA       | NA       | NA       | NA       | NA       | NA            |
| GOT 8               | -0.7     | -2.6     | -10.5    | 6        | Closed   | 419243477     |
| # show goat         |          |          |          |          |          |               |
| 11/20/2019 19:21:35 |          |          |          |          |          |               |
|                     | DSTX dBm | USRF1 dB | USRF3 dB | USRF4 dB | USRF6 dB | Serial Number |
| GOA                 | 10.0     | 0        | NA       | 0        | NA       | 519301888     |
|                     | DSRX dBm | USTX dBm | USRX dBm | USRF dB  | Node Lid | Serial Number |
| GOT 1               | NA       | NA       | NA       | NA       | NA       | NA            |
| GOT 2               | 0.2      | -2.6     | -9.0     | 0        | Closed   | 419243545     |
| GOT 3               | NA       | NA       | NA       | NA       | NA       | NA            |
| GOT 4               | -0.3     | -2.6     | -9.0     | 6        | Closed   | 419243566     |
| GOT 5               | NA       | NA       | NA       | NA       | NA       | NA            |
| GOT 6               | 0.2      | -2.6     | -8.6     | 0        | Closed   | 419243513     |
| GOT 7               | NA       | NA       | NA       | NA       | NA       | NA            |
| GOT 8               | -0.5     | -2.6     | -10.4    | 0        | Closed   | 419243477     |

**Figure 10 – CLI with Various Settings**

Figure 10 shows how these values can differ. Notably, the USRX values are nominally at -9 dBm because each of the US links have a ~6dB optical pad. However, since the return path is digital, the optical attenuation does not affect the RF levels at all. The only thing that affects the RF level is RF attenuation, which is shown here in Red. On the top, GOTs 2 and 8 are attenuated as is the GOA RF Block 1. In the bottom picture, only GOT4 is attenuated.

When GOT 4 output level is attenuated in the GOA, two things happen:

1. The noise and the signal level of the GOT4 node is attenuated by 6 dB. Therefore, if GOT4 was the ingress culprit, we would have ingress relief of 6 dB
  - a. Since all the RF was combined together, any reduction in ingress will beneficially improve the US SNR of the entire system
  - b. If the aim is to identify the ingress point, that is now accomplished. If the aim was to mitigate ingress, then the attenuation stays on until technicians can trouble shoot and eliminate the ingress
2. In addition, the long loop AGC administered by the combined vCMTS/RPD kick in and the RF output levels of all CMs connected to GOT 4 are increased by 6 dB
  - a. When this happens, the CMs connected to the GOT 4 benefit and their SNR improves as well
  - b. Some CMs that are on the high edge of USTX could reset. This is dictated by the duration of the attenuation and how close they are to the edge



**Figure 11 – Highly Simplified Ingress Illustration and the Role of US RF Attenuation**

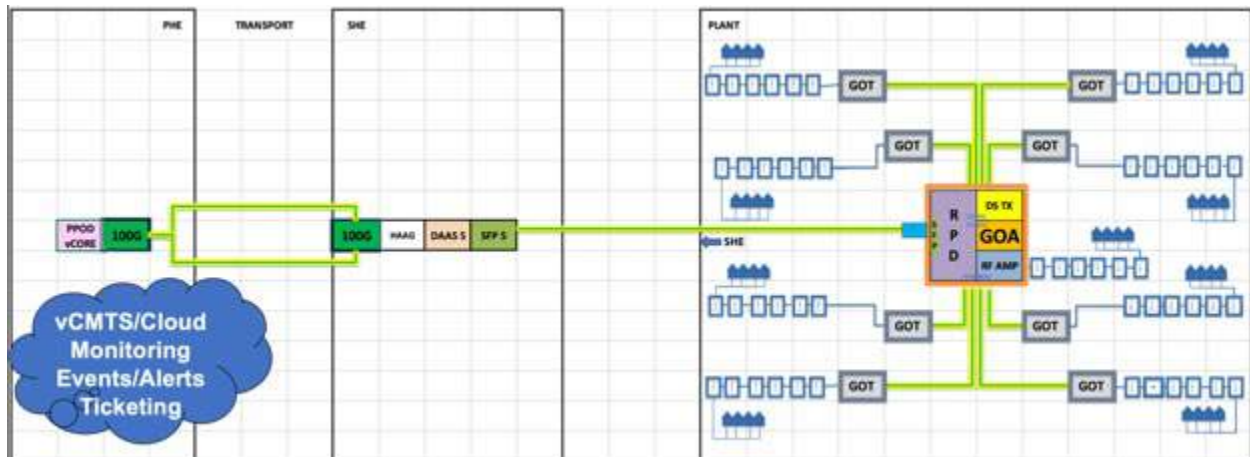
The above points are illustrated in Figure 11. Roughly 90% of upstream ingress accumulates at the home, as characterized in an excellent paper written by Larry Wolcott [5] who explains that minimal additions to it occur in the RF plant -- but the plant does funnel the ingress, which compromises the entire plant's SNR. At the extreme it is possible that an aggressive ingress source from a single household can impact the SNR of the entire system. Therefore, it is certainly possible that a single GOT can impact the SNR of the entire system. According to experts whose job it is to keep the upstream signal path functional, around 80% of trouble reports and technician time is spent on ingress identification and mitigation. Since ingress is generated at the home, it can occur in DAA plant and in GOA plant just as it can occur in a regular HFC plant. The ability to peer into the network remotely and identify the ingress-impacted GOT nodes or RF ports, and reduce its impact until it is all improved, is a potential game changer for the future of the RF plant. This capability, although developed for GOA, can be extended to any RPD node, whether it is deployed for DAA or HFC applications.

## 8. Tooling, Eventing, Layering ... All in a Sprint

Using a Command Line Interface (CLI) for quick checks is one thing, but using it to continuously monitor and control a large number of GOA RPDs is not a good idea. This is because repeated logging into the RPD (using a scheduling utility such as a Cron job) has a way of destabilizing the network. In addition, RPD logging is taken very seriously in Comcast and the security steps make the process much more difficult.



Fortunately, we have a set of software that uses TLVs (Type Length Values that are used in data communication protocols) that specify information of interest and “broadcast” these to those interested. Harnessing the TLVs by appropriately defining them and connecting them to the various pieces of information will then enable us to “listen in” on them, and create dashboards that enable us to look at various network elements in real time.




**Figure 12 – The Tooling Diagram – Cloud Based Provisioning/Monitoring/Tooling**

Within Comcast, a project or a product is never considered finished until the “tooling, eventing and layering” tasks are complete. This is a special lingo that takes a bit of time to absorb, particularly for a hardware optical engineer such as this author. But after a bit of learning, we focused on defining the TLVs for the various elements of the CLI table defined in Figure 12. Thus, all of this information would be continuously available (every 15s in the limit) and the entire constellation of the GOA, including all the GOTs and their US and DS levels and RF attenuation state.

This information is available to Grafana, an analytics and visualization tool, for read-only graphing and history purposes for anyone in the organization. This is especially useful to the engineers, and also to the technicians, division engineers/leaders and also to operations centers as necessary.

We also use a piece of software called the RPD Life Cycle Manager (RLCM). This is secure software that enables one not only to view all aspects of the RPD, but also to control aspects of it without additional security. Thus, the ability to manipulate the RF attenuation values are listed in the RLCM.

By now, it is clear that there is ample information available within the company to understand all aspects of our network. We continue to make strides in making all of this data actionable, with increasing application of machine learning and artificial intelligence algorithms. In fact, several of these aspects were used during the COVID pandemic to not only improve reliability but also increase capacity.



**Figure 13 – Visualizing All US and DS Metric at a Glance Across the Node**

For measuring performance of the network, however, we have so many metrics that it becomes difficult to look at all simultaneously, yet, looking at them is unsatisfying and probably dangerous as the limited information could lead us to incorrect remediation.

Consider a spider chart of the kind depicted in Figure 13. Here, 5 of the most important parameters have been coded to the table alongside and plotted. Since we have data on each modem in real time, these data are available for view on a modem by modem basis and available for display. It is easy to see here that for this node, the US RF RX is always close to ideal (0 dBmV in real life and 10 on the spider chart), but one of the CMs has a rather poor time of getting all other parameters out (DSRF MER is between 35-37 dB, DS RFRX is -8 to -10 dBmV, USRF SNR is 30-32 dB and the US RFTX is 51-54 dBmV). Meanwhile, all the other CMs are in much better shape. This allows us to know, at a glance, that this modem has a low DS RFRX and a high US RFTX value (probably due to a high tap value in an in-home network) which is why the DSRF SNR is low, so, improving it might serve all the parameters well. In general, the wider the “Spider Eye” is, the better the system is, and the eye opening may be used to verify overall performance in a non-invasive way as an added verification of the Orchestration process.

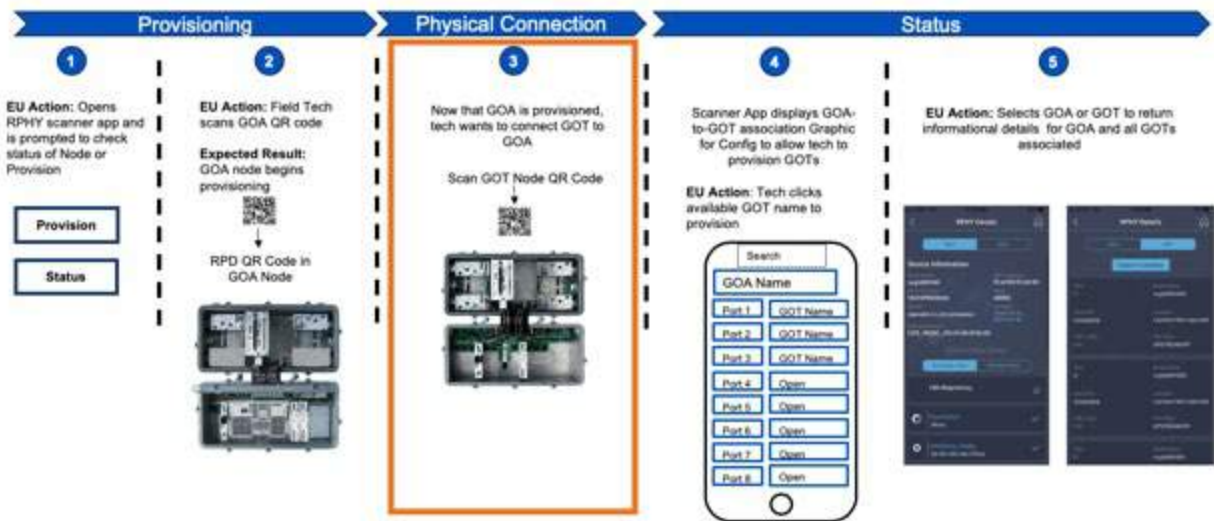
Together the RLCM and the RPD feed events to the internal National Watch Tower (NWT) engine. NWT is a powerful visualization and event reporting and ticketing tool that is one of our internal operational staples. For the architecture to function smoothly and be a template for the future crop of DAA RPDs, all these features must work seamlessly, and are next described in detail.

## 9. Provisioning the GOA-GOT system

We have now entered the Software World, where verbs function as nouns and all specific actions are gerund-sounding ... starting with the provisioning process.

Provisioning an RPD is usually done by scanning the QR code of the RPD and matching the IP address, MAC address and the node name together as a first step. The case is similar with a GOA RPD. However, for the GOT node, there is no MAC address, nor is there an IP address. We solved this lack of information by using a piece of software we developed called the RPHY Scanner App.

The RPHY Scanner App provisions an RPD and recognizes when the RPD is a GOA RPD. This can happen because the RPD is connected to a Receiver aggregating module and has a DS Transmitter in the node. Therefore, the App automatically opens up a GOT tab. Recall that we always go from the ‘Heart to Limbs’ for the GOT deployment. Since the GOA RPD is already live, when the fibers (of specified colors as described before) are connected, and the DS light enters the GOT, the US light leaves the GOT and enters the GOA node, the node is automatically given a name, which is a single digital extension of the GOA node name. For example, if the GOA were called NGACD00110, then the GOT connected to GOA receiver 1 would just be NGACD001101, and the one connected to receiver 2 would be NGACD001102, and so on.



**Figure 14 – Provisioning Application Flow**

With a name available before the GOT is provisioned, the technician is prompted to take a picture of the GOT QR code. When that is done, the serial number of the QR code is matched against the serial number that is sent via the TLV. Assuming they both match, the Latitude/Longitude of the GOT position is stored and the node is declared provisioned. Thus, each GOT node is provisioned, its Lat/Long position set, its serial numbers verified, and it can start as a fully functional node in the GOA constellation.

## 10. “Grafana-ing” the Constellation

Grafana is a formidable tool that enables one to graph the performance of any device over time. As discussed, we get a range of US and DS RX and TX information along with data on the RF attenuations, in real time, every 15s or so. It should therefore be possible to slice that information and present it in a graphical form. Such a task is achieved in the ‘GOA Constellation’ part of Grafana as a read-only feature. This is especially helpful to see trends, identify intermittent issues, or perform historical comparisons (such as day to day or week to week).



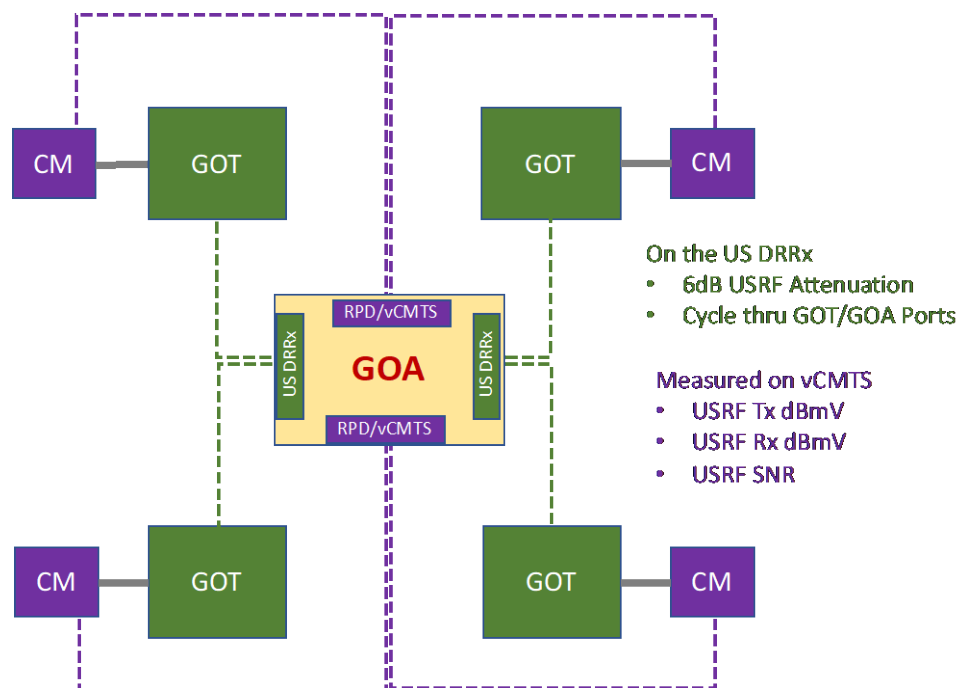
**Figure 15 – Grafana and Real Time and Historical Constellation Information**

In Figure 15, the following can be seen:

1. Each GOA can be individually selected for viewing based on the PHE location or the node name. The timeline can be varied as well.
2. The top portion of the dashboard shows the DSTX level in real time as a graph. The value is nominally 10 dBm, as we have seen before. Each of the DSRXs at each GOT are shown as well, and we know that the nominal values are -0.5 dBm. If, for any reason, the DS fiber were cut/disconnected, this would show as a disconnect and point to the issue.
3. The middle of the dashboard shows the USTX and the USRX data for each GOT. Here it is seen that the nominal value of the DSTX is ~-2.5 dBm and the nominal value of the USRX is ~-4 dBm, as can be seen in the drop-down list.
4. In the bottom part of the dashboard, one can see the RF attenuation experienced by the various ports. Recall that the GOA has 2 RF Ports, and each of the GOTs is treated as an RF port for purposes of attenuation, to identify and mitigate ingress. Here one can see the 8 GOTs undergoing RF attenuation, one after the other. Even though the attenuation is applied for a uniform time, the graph is limited by the resolution of the tool and the display and the time it takes for all the information to be collected.

This tool has been of utmost importance, as we have used it extensively to study the robustness of our constellation, especially as COVID raged across our country, with its unquestionable impacts on available capacity. Understanding these issues was impossible with direct observation. The vast amount of historical data is a perfect complement to the CM data and traffic data that is generally available and stored in other portions of the Grafana tool. This portion of tooling is vastly complex and not discussed further.

## 11. Orchestration and Association



**Figure 16 – Illustrating the concept of Association and Orchestration**

It is common for the CMTS to know all the CMs connected to it, as well as the modem RF transmit power, the modem RF receive power, and also the RF receive power at the CMTS input. In this case, the vCMTS is expected to keep track of these parameters in concert with the RPD (since features of the CMTS have been distributed between these two by principles of virtualization described in the beginning). We have further seen that the GOA node knows about the existence and state of all the connected GOT nodes, and is aware of the optical DS input power, the optical US transmit power and the US received power. But since the US is a DRT/DRRx connection, changes in optical power do not affect RF power received power as long as the link is functional. Of course, we would not like the optical power to change because this will cause events. We will discuss how these are flagged and ticketed in Section 12.

In Figure 16, the situation described above is represented by the purple dashed lines connecting the RPD/vCMTS to the CMs, to signify that the US RF TX and US RF RX power, as well as the US RF SNR is known. In addition, the DS RF SNR and the DS RF MER is known. The green dashed line represents the connection between the GOA and GOTs via the constellation program described above. What we do not have is the knowledge of how the CMs are connected to the individual GOTs. Nor do we know if ingress exists, and if so, where it might be originating. Note here that even though the GOA itself is also connected to the plant, we have not depicted it, as a way of simplification.

The ability to map the individual CMs to the individual GOTs is called Association. To do this, we orchestrate 6 dB RF attenuation one after another on the various GOTs in the DRRx in the GOA node. This way, the RF levels of the CMs that are connected to the specific GOTs are raised by 6 dB, due to the long loop AGC of the CMTS in response to the attenuation. By keeping track of the US RFTX, US RFRX and the US RFSNR, we can pinpoint the CMs associated with individual GOTs, and also, by tracking the SNR improvements, we can pinpoint the magnitude and location of any ingress sources.

By way of explanation, we keep track of the baseline of all US RFTX, US RFRX and USRFSNR for each of the CMs connected to the GOA RPD. We then put in 6 dB attenuation on one of the GOAs, which will cause the RF levels of all CMs to rise by 6dB. Furthermore, we track if the SNR of the system has improved, and the extent of any improvement. All CMs that reacted by increasing their RF levels are associated with the GOT node. If the SNR improved, then said GOT is identified as an ingressor. We then rate this on a scale of 1-6, with 6 dB being the most egregious ingressor.

With proper verification, we have been able to secure 100% accuracy of the CM associations, in real time, remotely. Because DAA deployments tend to have RF transmit variables well within limits, we have seen very few (~1%) of CMs reset with changes in RF levels, and they come back online within the duration of the test, once the attenuation has moved away. We have also been able to track ingress events remotely, which was immensely helpful in the COVID environment. This is because technicians can approach the exact node and troubleshoot for ingress in the rare events that presented themselves during the deployment. This reduces their time in the field and reduces any associated virus exposure in customer homes. How exactly this is implemented securely, remotely and automatically is described next.

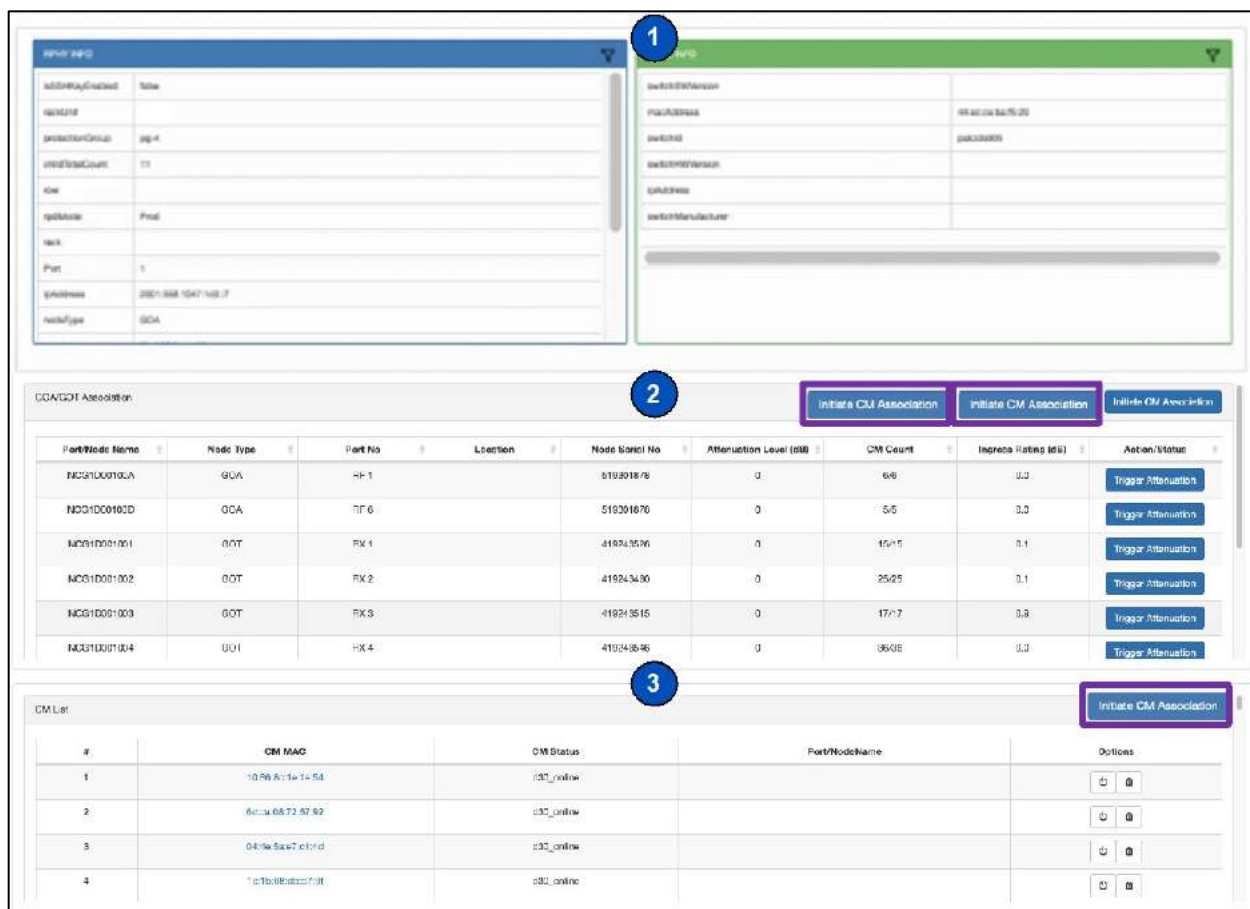
## **12. Secure, Remote and Automatic: RPD Life Cycle Manager**

Enabling DAA is a complex operation. It not only required the vCMTS scheduling effort, but a similarly complex back office system to monitor and control various parameters that weren't normally visible or useable in real time. We described the provisioning application, as well as the traffic and connectivity managing application; next we described the process of using a complicated real time learning app that helps the RPD monitor and control the various parameters of devices connected to the RPD. We have previously indicated that the use of the I2C bus in the GOA node allowed us to be able to surface all aspects of the node to the RPD monitoring system. This system already exists with RPDs and has been updated to include the GOA-GOT system.

The RPD is controlled remotely via the RPD Life Cycle Manager Application (RLCM). This application identifies the RPD as a GOA. This then opens up a table with the various GOTs and their serial numbers and their latitude and Longitude information (recall that this is secured when the GOTs are provisioned). The App further maintains a list of all current CMs, and their state, as is expected of a CMTS. This is based on direct observation by the RPD/vCMTS and not based on inference of any sort.

The App has a soft button for association, which, when pressed orchestrates the 6dB attenuation periodically, using machine/logic in the cloud, and establishes the CM association and the ingress rating. A button allows one to download the data for further analysis. Currently this is an individual operation. The thought is to expand it for periodic data harvesting, and to use machine learning to understand the nature, frequency and magnitude of ingress. This will provide a lasting improvement to any cable system.



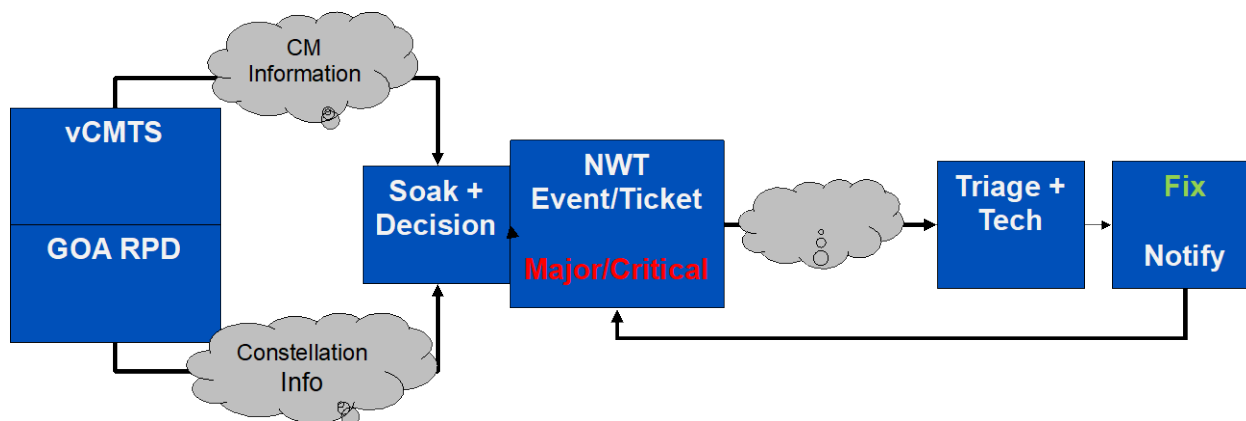


**Figure 17 – Association and Orchestration in the GOA RPD**

## 13. Eventing and Ticketing

The optical connectivity between the GOA and the GOTs is normally quite reliable, given that the fiber connections are typically less than a km. However, due to natural disasters and other events it is possible that the connectivity would get disturbed. If the DS and/or US fiber is cut, or if, for some reason, the GOT should lose connectivity to the GOA, it is imperative to identify the impacted GOA and GOT with a ticket to flag them for remedial actions.

It is, however, the case that information reaches the decision engine to call an event in multiple forms, all with delays peculiar to their content. If the fiber between the GOA and the GOT gets cut, the CMs go offline and this information reaches the decision engine via the vCMTS path. The fiber cut will also result in the GOA US RRX missing light on a previously lit receiver, which will then propagate this information to the decision engine. There may be outages attributable to power which could take down multiple nodes; this information reaches the decision engine via the powering network. Since outage information from various sources is arriving with various minor delays, if the decision engine begins reacting immediately, it will have reacted with insufficient information. Thus, to avoid this “race” condition, the decision engine “soaks” or stores all incoming information with various minor delays to look at it holistically.



**Figure 18 – Automatic Eventing and Ticketing**

Having complete information from all sources, by inference and direct observation, the engine declares an informational/major /critical event, and an application such as National Watch Tower generates a ticket and sends it to the concerned division/headend/operations center. At the center, it is triaged and allocated to the technicians. All GOA-GOT outages are identified by a GOA and GOT node name and also the Lat/Long locations of both devices. This is to give as complete information as possible, to not only identify the devices, but also point out the likely fiber route that may have been taken. In the NWT, the GOAs and GOTs are layered, so that it becomes easy to match the outage to the map area.

Once the technician has fixed the problem, whether fiber cut or node maintenance or ingress event, the event is cleared by the technicians themselves. This way the node is brought back to an event-free state and normal operation resumes.

## 14. Future Steps

This year, we reported on the lessons learned during Comcast GOA trials and deployments. In addition to proving out the technology and cost benefits in the field and observing its stability, deploying the architecture also enabled us to interconnect and create an eco-system of remote monitoring and control of the various GOA elements. We will continue to report on our ability to continuously monitor Grey Optical Terminating (GOT) nodes, detect cable modems connected to specific GOT nodes and remotely identify and mitigate ingress in the GOA domain.

Many of the remote techniques were used during the COVID pandemic and can potentially be generalized further to enable remote monitoring and mitigating capabilities across the network. Furthermore, in the future we will report on recent work on bi-directional Coherent Optical modules that enable efficient extension of the architecture that, over a period of time, can create cost effective convergent and scalable ‘Switch on a Pole’ access networks.

## 15. Acknowledgements

It is our pleasure to acknowledge the entire team within Comcast who has been working directly and indirectly on the GOA project. A project of this magnitude benefits from the dedication of diverse expertise, from the hardware and software of the RPD team, the reliability and functionality testing of the Physical and Environmental team, the ticketing and alerting of the NWT team and the functional rules for deployment from the access engineering team for both ISP and OSP deployments. Special thanks to our



vendor partner CommScope for your insights. We sincerely thank the Senior Leadership Team at Comcast NGAN in supporting this project.

## Abbreviations

|         |                                            |
|---------|--------------------------------------------|
| AGC     | Automatic gain control                     |
| CLI     | Command Line Interface                     |
| CM      | Cable modem                                |
| CMTS    | Cable modem termination system             |
| DAA     | Distributed access architecture            |
| DAAS    | Distributed access architecture switch     |
| DRT     | Digital Return Transmitter                 |
| DS      | Downstream                                 |
| DSRX    | Downstream receiver                        |
| DSTX    | Downstream transmitter                     |
| EMS     | Element Management System                  |
| FEC     | Forward error correction                   |
| GOA     | Grey optics aggregation                    |
| GOT     | Grey optics termination                    |
| HHP     | Households passed                          |
| HSD     | High speed data                            |
| ICS     | Ingress control switch                     |
| MER     | Modulation Error Ratio                     |
| NGAN    | Next Generation Access Network             |
| PER/FEC | Packet Error Rate/Forward Error Correction |
| RF      | Radio frequency                            |
| RFSNR   | RF Signal to Noise Ratio                   |
| RLCM    | RPD Life Cycle Manager                     |
| RPD     | Remote PHY device                          |
| SOAP    | Switch on a pole                           |
| SFP     | Small form-factor pluggable                |
| SNR     | Signal to noise ratio                      |
| TLV     | Type length value                          |
| US      | Upstream                                   |
| UPRX    | Upstream receiver                          |
| UPTX    | Upstream transmitter                       |
| vCMTS   | Virtual cable modem termination system     |

## Bibliography & References

1. *Fifty Shades of Grey Optics: A Roadmap for Next Generation Access Networks*, Venk Mutalik, Bob Gaydos, Dan Rice, Doug Combs, SCTE Expo 2019
2. *Distributed Access Architecture – Goals and Methods of Virtualizing Cable Access*, Nagesh Nandiraju et. al., SCTE EXPO 2016

3. *When Wavelengths Collide, Chaos Ensues: Engineering Stable and Robust Full Spectrum Multi-wavelength HFC Networks*, Venk Mutalik et. al., SCTE Cable-TEC EXPO 2011
4. *Cable's Success is in its DNA: Designing Next Generation Fiber Deep Networks with Distributed Node Architecture*, Venk Mutalik et. al., SCTE EXPO 2016
5. *A Comprehensive Case Study of Proactive Network Maintenance*, Larry Wolcott, John Heslip, Bryan Thomas and Robert Gonsalves, SCTE EXPO 2016

# **It's 10 PM: Do You Know Where Your Wavelengths Are?**

## **Continuous and Pervasive Monitoring of Optical Assets in the Access Domain**

A Technical Paper prepared for SCTE•ISBE by

**Venk Mutalik**

Executive Director, HFC Architecture  
Comcast  
1401 Wynkoop St, Denver, CO  
+1 (860)-262-4479  
Venk\_Mutalik@Comcast.com

**Dan Rice, Comcast**

1401 Wynkoop St, Denver, CO  
Daniel\_Rice4@Comcast.com

**Rick Spanbauer, Comcast**

1401 Wynkoop St, Denver, CO  
Richard\_Spanbauer@Comcast.com

**Simone Capuano, Comcast**

1401 Wynkoop St, Denver, CO  
Simone\_Capuano@Comcast.com

**Rob Gonsalves, Comcast**

1401 Wynkoop St, Denver, CO  
Robert\_Gonsalves@Comcast.com

**Bob Gaydos, Comcast**

1800 Arch St, Philadelphia, PA  
Robert\_Gaydos@Comcast.com

# Table of Contents

| Title                                                           | Page Number |
|-----------------------------------------------------------------|-------------|
| 1. Abstract .....                                               | 3           |
| 2. Introduction .....                                           | 3           |
| 3. Access Architecture .....                                    | 5           |
| 4. Fiber Connectivity and Inference .....                       | 7           |
| 5. Current State of the Art .....                               | 8           |
| 6. Pervasive Monitoring Paradigm .....                          | 8           |
| 7. Role of Optical Passives .....                               | 12          |
| 8. Infrastructure Evolution .....                               | 13          |
| 9. End-to-End Optical Architecture .....                        | 16          |
| 10. Praemonitus est Praemunitus (Forewarned is Forearmed) ..... | 18          |
| 11. Provisioning and Training the Monitor .....                 | 19          |
| 12. Handheld to the Cloud .....                                 | 22          |
| 13. Conclusions .....                                           | 23          |
| 14. Acknowledgements .....                                      | 23          |
| Abbreviations .....                                             | 23          |
| Bibliography & References .....                                 | 24          |

## List of Figures

| Title                                                                               | Page Number |
|-------------------------------------------------------------------------------------|-------------|
| Figure 1 – Comcast Long Haul Fiber Network .....                                    | 4           |
| Figure 2 – Comcast Residential Network Areas .....                                  | 4           |
| Figure 3 – Simplified Connectivity Diagram .....                                    | 5           |
| Figure 4 – End to end DAA Architecture Diagram .....                                | 6           |
| Figure 5 – Highly Simplified view of the Continuous Pervasive Monitoring Tool ..... | 10          |
| Figure 6 –Micro and Macro Bending Losses for ITU Rec G.652 Fiber .....              | 11          |
| Figure 7 – Optical Passives Old and New in Comcast .....                            | 13          |
| Figure 8 – Actual Comcast Installation of the Continuous Pervasive Monitor .....    | 14          |
| Figure 9 – Describing the DAAS Pod and the Monitor .....                            | 15          |
| Figure 10 – End-to-End Software Paradigm .....                                      | 16          |
| Figure 11 – Sample view of the OSA and OTDR Output for each port .....              | 17          |
| Figure 12 – Sample view of the OSA and OTDR Output for each port .....              | 18          |
| Figure 13 – View of the Provisioning page in for the Monitor .....                  | 20          |
| Figure 14 – Fully Provisioned Port (notice the green and red WL colors) .....       | 21          |
| Figure 15 – Handheld unit with a Mobile Application and Cloud connectivity .....    | 22          |

## 1. Abstract

Today, the access fiber plant carries large amounts of varied data content deeper into the network. Analog wavelengths support traditional Hybrid Fiber-Coax (HFC) plant, 10 Gbps wavelengths support Distributed Access Architecture (DAA) network segments and other MetroE applications, while newer Coherent Optics modules support “Switch on a Pole” (SOAP) [1] type architectures. As use of the fiber footprint grows, so too will the need for comprehensive monitoring, to optimize the efficiency of access optical assets with the ability to inventory bi-directional wavelengths. These optimizations aid with capacity planning and to locate fiber cuts and other impairments across the plant, in real time, while identifying effective mitigation options.

For too long, such monitoring has solely been the purview of long-haul networks, but recent innovations in optical technology pioneered in part at Comcast enable automatic, continuous and pervasive monitoring of access optical assets without active user intervention. These cost-effective, switched optical devices, comprising an optical spectrum analyzer and optical time domain reflectometers modules, are co-located with access headend optics and continuously scan links to detect fiber cuts or individual wavelength outages. These headend tools are augmented with the same cost-effective optical measurement technology implemented in handheld meters for field testing. Headend and handheld tools together provide real time optical data to the cloud for data-analytics. Locating an access fiber-cut in real time, and automatically guiding the Comcast response team to its exact location, results in exceptional uptime, which enhances customer satisfaction -- especially in periods of disaster recovery.

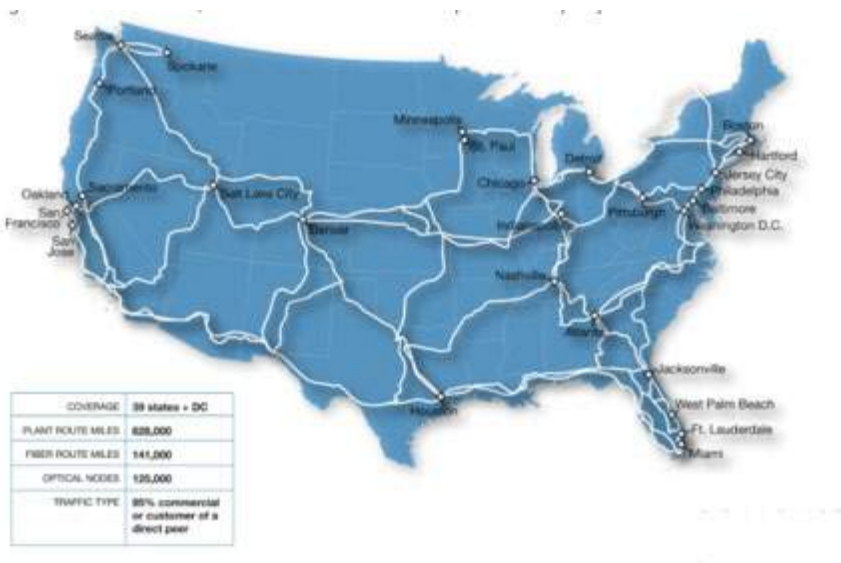
In this paper, we begin with a description of the optical innovations and powerful techniques that enable this extraordinary tool set. We then describe the infrastructure that was stood up to intake, visualize and “event” this data, and detail our preliminary experience using this technology in the COVID lock down period. We will then venture into the future of continuous pervasive monitoring of access optical assets, and the positive impact such next-gen monitoring has on network robustness, which, in turn, enhances the customer experience.

## 2. Introduction

Readers will perhaps remember the popular public service announcement "It's 10 p.m. Do you know where your children are?" intoned before the nightly news. This was memorably used in a 1996 episode of *The Simpsons*, when Homer Simpson responds to his TV, "I told you last night — no!" And so it was with our fiber plant for many initial years -- as the plant grew in density and reach, then added multiple wavelengths for each fiber, we were often in Homer's position, when it came to knowing where our wavelengths were, at 10PM or otherwise!

With improvements in technology and design, and as our long haul and back bone networks increased capacity and reach, we initiated programs to track and keep knowledge of wavelengths used, fiber cuts and impairments in the plant. With the advent of Remote Optical Add-Drop Multiplexers (ROADMs) and more sophisticated technology, this knowledge was now available, and in real time.

Comcast is one of the largest broadband providers in the country. Our network reaches from coast to coast. In total, in the U.S., our network comprises ~150,000 miles of fiber route miles, as shown in Figure 1.



**Figure 1 – Comcast Long Haul Fiber Network**

Part of running a high-capacity backbone network is the ability to monitor and troubleshoot any outages and disruptions due to fiber cuts and/or equipment malfunctions. The larger the network, the more challenging this task.

Similarly, we run a very large footprint of Access Plant, covering ~55 Million Households Passed (HHPs), and illustrated in Figure 2. This access plant network (marked in dense Red points), also called Hybrid Fiber-Coax (HFC), generally has Cable termination to the end customers, but a vast majority of the access plant is also covered by Fiber, from the Primary headend (PHE) to the Secondary Headends/Hubs (SHE) to Fiber Nodes.




**Figure 2 – Comcast Residential Network Areas**

With new Distributed Access Architectures (DAAs), we anticipate the fiber content of this network to grow exponentially over time -- and by over 10x in the coming few years. With this huge growth in our fiber assets, it is now imperative to more tightly monitor and secure our high-capacity access networks as well.

In the remainder of this paper, we describe our efforts to create a hardware and software infrastructure that continuously and pervasively monitors our optical assets. We also describe how such monitoring can come in very handy during the present COVID environment, and likely in perpetuity.

### 3. Access Architecture

Figure 3 presents a highly simplified connectivity model for our access architecture. In big markets, we typically have a high availability fiber back bone linking the Primary Headends (PHEs). These are typically connected to Secondary Headends (SHEs) via a redundant route. Often the SHEs themselves are arranged in a ring connected off of the PHEs, also with redundant fibers, which then connect to individual fiber nodes.

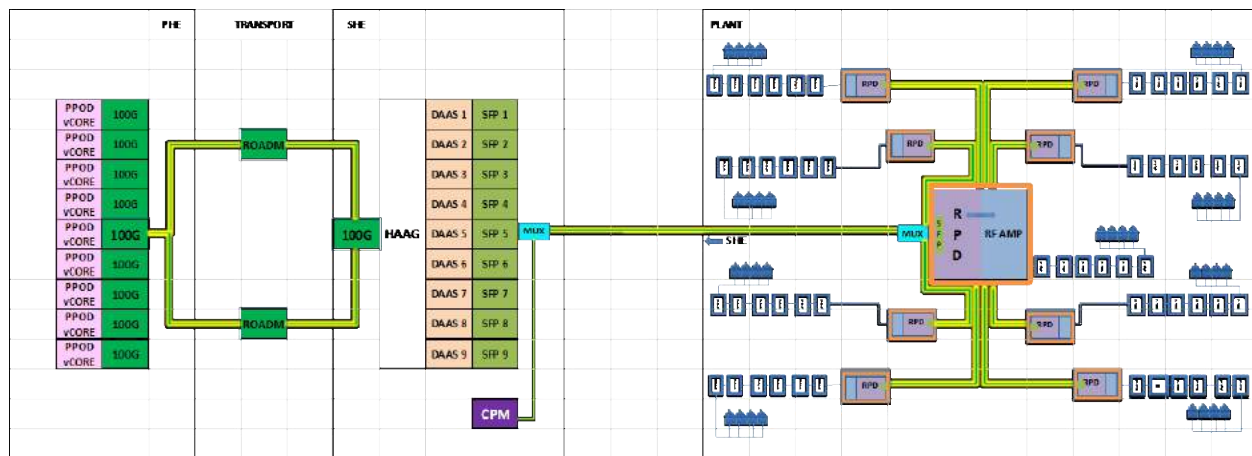


**Figure 3 – Simplified Connectivity Diagram**

Typically, the PHE network is supplied by high speed (such as coherent 100G-400G) links which have levels of monitoring for all optical wavelengths in use, and capabilities for locating fiber cuts or impairments.

As can be seen, for each PHE, there are several SHEs -- and for each SHE, there are a lot of nodes. While the SHEs are connected directly to the PHE with redundant routes, the nodes are directly connected to the SHE with straight fiber connections. Thus, the access plant has a lot more optical connections. Even considering the coaxial plant exiting the fiber nodes, the sheer amount of fiber links in the access plant is much greater than what exists in the long-haul plant.

With the Distributed Access Architecture (DAA), each of the parent nodes is now connected via fibers to multiple nodes within its footprint. We anticipate that the fiber plant will grow exponentially in the access domain.



**Figure 4 – End to end DAA Architecture Diagram**

In the DAA architecture shown in Figure 4, PHEs would contain all of the virtualized CMTS (vCMTS) cores, connected to the SHEs through redundant routes to DAA Switches, and then via SFPs to the individual Remote PHY Device (RPD) nodes. These RPD nodes are connected to individual homes across coax. The return circuit is connected back via the RPD node to DAA Switch (DAAS) and then through to the vCMTS in the PHE.

Several other vCMTS-based architectures are in the process of being deployed and worthy of mention here. In one instance, the RPD is itself within the SHE and an analog link connects to the fiber node in the field. In this case, the trunk fiber would use Dense Wave Division Multiplex (DWDM) wavelengths on the 4 Wave Mixing mitigated plan (aka the “Full Spectrum” [FS] plan.) The reverse path typically uses baseband digitized return links (BDR links). In this case, the monitoring requirements would not change, but rather move to monitoring the analog links that have thus been established.

In other variations, the RPD node would still be in the field, but would encompass additional passings. This could be because of an internal split (such as 1x2 or 2x4), or due to RF amplifiers that exist in the plant. In this case, we have a need to monitor the digital links as described earlier.



In yet other variations, we may have analog nodes deeper into the plant with no RF amplifiers. These would have the downstream (DS) on analog FS plan described earlier but would have separate fibers for the upstream. What complicates the situation further is that these links have asymmetric US and DS wavelengths per fiber, thus resulting in asymmetric DS and US fiber counts. And the need is still to monitor these fiber links with a view to know individual node connectivity and inventory fiber capacity.

The vast majority of nodes, though, are the regular HFC nodes that are connected to the headend with either a single WL, or with multiple WLs on the FS plan, with US WLs sometimes on the same fiber, and sometimes on other fibers. All in the need of fiber monitoring for connectivity.

It is easy to see here that the sheer number and the variety of links present a unique opportunity for operators like us to improve further on our network availability to the node and prepare for the future. The difficulty has thus far been the cost of getting said monitoring developed and in place, plus the intensive manual nature of troubleshooting, which is discussed next.

## **4. Fiber Connectivity and Inference**

An outage, particularly a suspected fiber cut, instantly calls attention to a wide variety of resources in Comcast. With advanced internal tools like National Watchtower (NWT), we are able to know within minutes that individual Cable Modems have lost touch with the CMTS. The inference engine in the NWT can identify a problem node, to which trouble tickets are then generated for workforce management. If the whole node is out, then chances are high that a fiber may have been cut, heavily impaired or involved in a utility power outage. An optical time domain reflectometer (OTDR) is pressed into service to determine fiber connectivity, after a quick check that the return light is absent at the headend. The OTDR reveals a fiber cut, which is then located using a Geographical Information System (GIS) and fiber plant data. Then, a crew is sent to the spot for repair. When they get there, the restoration duration depends upon the number of fibers to be re-spliced.

In such cases, it is quite possible that the fiber sheath that is cut could also house a number of fibers, some of which could serve multiple non-residential services, such as Metro Ethernet (Metro E) or other business services, all managed by Service Level Agreements (SLAs.) Sometimes, especially if the fiber cut location is inconclusive, field technicians have no choice but to OTDR the fiber backwards, to further attempt to locate the cut. This adds both time and frustration. After the cut is restored, fiber plant is put back and the ticket cleared.

If, however, the node is affected, but it is not due to a fiber cut (which is known after an OTDR shot is taken), or if other nodes known apriori on the same fiber are unaffected, then, without a power outage, the dispatch instructions for the node would be to replace its optics/RF components. All of this information is available in multiple places, and the amount of time effort and energy it takes to communicate effectively when responding to such incidences continues to improve, steadily, as tooling and integration improve. However, this is still a logical step in addressing the above need and improve the customer experience further.

## 5. Current State of the Art

Although technology has improved in major ways, and costs have come down, the current state of the art in mobile optical spectrum analysis and fiber connectivity is still quite expensive and localized. It is not unusual to see major equipment at headends and equipment moved around as needed. In an effort to lower cost, optical channel checkers have proliferated that enable us to measure individual WL power values. But these do not have integrated OTDRs. Many in the optical equipment manufacturing community now offer modular OSA and OTDR models -- but adding too many modules make the equipment bulky, and having fewer slots requires module swaps, which can be an issue in the field.

Since most of our access network footprint is in the 1550 nm region, using a standard 1550nm OTDR would not enable live scoping of the fiber plant since the OTDR light pulses could interfere with the smooth functioning of signals if the wavelengths coincide. Thus, this becomes used only in known or suspected fiber cuts or impairments. Mobile tunable 1550 nm OTDRs are a great innovation and could be used in live plant on unused channels, but we would need to apriori know that those channels are not in use. In addition, care must be taken in their use, to avoid damaging connected SFPs at the far end, because of damage or impairment to Avalanche Photodiode (APD) receivers in SFPs.

While channel checkers built with a filter/photodiode combination would track power levels, they will not be able to track wavelength movements, which can sometimes occur in the plant. Furthermore, it is not uncommon to use non-ITU wavelengths in the plant, such as when a broadcast transmitter is deployed in an overlay system. Comcast developed an architecture option for grey optics aggregation (GOA) that, as the name implies, uses Grey Optical Transmitters, where the transmitters are in the 1550 nm region without them being on any one specific ITU wavelength. In such cases a channel checker is insufficient and an OSA function is needed.

In a different paper [3], we discuss a more general convergence model at Comcast, where multiple services are all combined together on the same single fiber bidirectionally. In that model, analog DS WLs along with BDR signals of varying speeds, 10G RPD node connections, 10G Metro E connections, and 100G connections are all multiplexed on the same fiber.

## 6. Pervasive Monitoring Paradigm

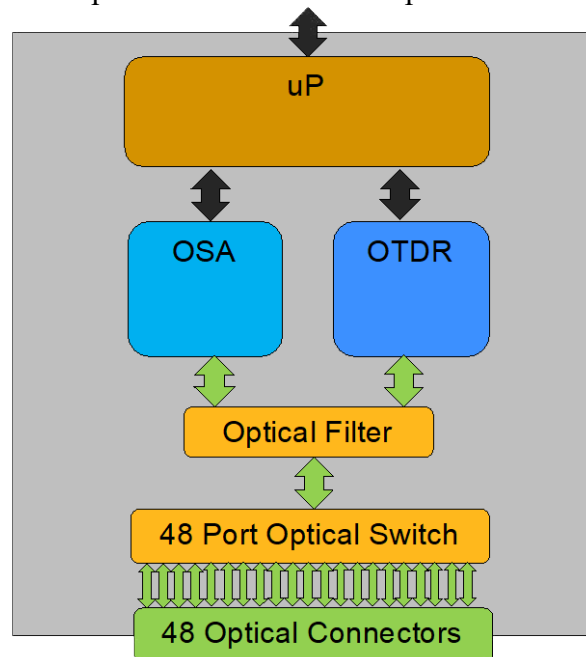
Our thought process for continuous and pervasive monitoring began with three fundamental insights:

1. Use of commercially available components to cost effectively monitor individual wavelengths and to detect fiber cuts across the access plant
2. Continuous monitoring of all optical assets, using only test ports, without impacting the fiber reach of the access plant
3. Pervasive monitoring of optical assets in a cloud-based environment is essential to efficient correlation and problem resolution (as opposed to targeted stand-alone solutions.)

To elaborate further, it is observed that OSA chips are widely used in ROADM applications. Cost effective and plentifully available, these chips used in ROADMs in the test port path could enable us to view optical wavelengths in real time. Many of these chips have a ~500ms scan time for the entire C-Band, and function from ITU 14 thru ITU 62 range, which makes them capable of detecting and showing WLs regardless of being on the ITU spectrum.

Selecting an OTDR of 1611 nm enables cost effective fiber connectivity information over live plant. Furthermore, 1611 nm Coarse Wavelength Division Multiplexing (CWDM) optical passives are cost effective and plentifully available. Additionally, single mode optical fiber is weakly guided and much more sensitive to micro and macro bending losses at 1610 nm (than it is at 1550 nm), so that “fiber choking” described earlier is seen much earlier than on the signal. 1310 nm is much stronger guide and has higher losses over fiber, thus is less suitable than 1611 nm for this purpose. We also selected 1611nm vs. 1620 nm due to the more widespread availability of optical passives at this region, and the fact that we generally do not use the 1611 nm window in the access plant (except in FTTH/RFoG deployments). We have already indicated that a standard 1550 nm prevents live channel OTDR and violated our rules of pervasive and continuous use, thus deeming it unsuitable. A DWDM Tunable OTDR would be much more expensive and require spare ports, and therefore that option was also not selected at this time.

Even with cost effective and capable OSA and OTDR options, we still required additional insight into making this technology pervasive in general operating conditions. That insight led to the use of an optical switch, to enable a “round robin” perpetual scan of the multiple fibers in the access plant. As described earlier, the fibers from SHE to the node are multiplexed at the SHE and demultiplexed at a field location, then feed individual nodes. Several such links exist in the SHE -- typically one may see up to 128 fiber links in each SHE. Therefore, the ability to have a 48-port switch that can connect to the test ports enables us to cost effectively deploy the OSA and OTDR functionality over up to 48 links. This concept is shown in Figure 5.




### Figure 5 – Highly Simplified view of the Continuous Pervasive Monitoring Tool

As an aside, it is common to see fiber counts in multiples of 12 as opposed to in powers of 2, perhaps because 12 has many factors that make it convenient to divide a fiber's route after the fact. In any case, we extended this number of wavelengths available on optical passives to up to 48. Therefore, it is possible that one OSA, OTDR combination behind the 48-port optical switch to survey  $48 \times 48 = 2304$  WLs, which can translate to 1152 optical WL pairs or nodes.

With this background, a pervasive monitoring tool was built, as described next. We began by selecting an appropriate OSA chip currently used in ROADMs with sufficient resolution (say 6.25 GHz). Typically, these are built using temperature-sensitive elements with a calibrated receiver that has a well-established relationship between temperature, wavelength passed and photodiode responsivity. This way, the whole C-Band is scanned, and power at each step classified within ~500ms. Tying up the OSA component to ROADMs gives us a reliable and sustainable way to ride their wave of cost effectiveness and performance enhancements. OTDRs function by sending out a series of pulses, receiving echoes, then analyzing the delay in the echoes to estimate the points of reflections. There is a fair amount of ambient reflection in fiber, due to Rayleigh scattering, and OTDR sensitivity is typically improved by higher laser power. The higher the laser power, the more the reach and the better the echo received back that enables one to estimate reflections. Sensitivity also improves with using an APD receiver vs. a PIN receiver. Finally, care must be taken to reduce any immediate initial attenuation, and therefore it enhances the launch power into the fiber and consequently improves sensitivity. It must be noted that all OTDRs have an initial dead zone that prevents identification of specific fault locations in close proximity to the OTDR itself. In our case we typically set fiber monitoring after about 250 m of the OTDR, which comfortably exceeds our dead zone.

Figure 6, reproduced with permission from Corning, shows the relationship between wavelength of operation and its effect on micro and macro bending losses for standard G.652 fiber (lower macro and micro bending losses are achievable using the ITU G.652D compliant fiber). Notice that the macro bending losses are low when the wavelengths are below 1400 nm due to exceptionally well guided light there. Also, notice how the losses are higher at 1550nm, but the losses increase dramatically past the 1550nm area. The micro bending losses also increase with wavelength, but the delta is rather limited.



**Figure 6 –Micro and Macro Bending Losses for ITU Rec G.652 Fiber**

A typical New England winter brings with it a freezing cold followed by a thawing cycle repeated over several months. Over time, many of the splice boxes get filled with water and “choke” the fiber when frozen. This can attenuate fiber significantly, sometimes disabling a link. Unfortunately, this usually happens at inopportune times, the weather being bad and all, and resolving this issue has been a bugbear. Fortunately, the selection of 1611 nm enables us to see the drop in light levels as an event far earlier than we would have seen on the signal itself. This is because the fiber weakly guides 1610 nm and more strongly guides 1550 nm. In fact, it guides 1310 nm much more strongly, which is why it is a poor candidate for OTDRs wavelengths -- besides the fact that the fiber loss is much more at 1310 nm, and for that reason also has poorer sensitivity and reach.) Higher sensitivity in the OTDR not only enables one to scan longer reaches, but also to make out individual losses in the fiber path, attributable to bad splices or micro- and macro-bending.

The time required for OTDRs to declare a fiber length is typically dependent on pulse widths, anticipated fiber lengths and the number of averages needed to get a clear idea of the fiber distance. This computationally intense process could easily exceed one minute and sometimes takes several minutes. Since we need to be able to scan in round robin fashion through 48 individual ports, the dwell time in each port would need to be multiplied by 48 in order to calculate the maximum time before a fiber break is identified (for sticklers of accuracy, one must also take into account the switch time, which is ignored for the moment.) Even a 1-minute dwell time would require 48 minutes before a fiber cut could be declared. In fact, this is a major difficulty that limits very drastically the number of ports realistically used.

For the pervasive monitor, we set up a target of declaring a fiber cut within 3 minutes for any one of the 48 ports. Our effort began with a commercially available OTDR module that offered the right set of optical specifications for our application but required more time than we had budgeted for our polled-scan method of monitoring the optical links. After some consideration of the OTDR requirements, it was clear that moving some of the digital signal processing from the original FPGA-based soft-CPU, to the high performance ARM used in the CPM, and increasing the data transfer rates at a key point in the OTDR module, yielded significant speed-up in the scan time. Only minor rework of the OTDR hardware was needed to implement this. Systematically understanding the computational requirements, the time of exporting the data for calculations, and then to convey the computed result, along with any events (more about that in a bit) enabled us to reduce the dwell time to just under 3 seconds! Furthermore, the OSA and the OTR could run simultaneously and independently on live plant. Thus, we are able declare any wavelength outage or fiber cuts within a maximum of 3 minutes across the plant. This massive and substantial improvement truly enables us to understand our plant in real time.

As described above, enabling the OSA and the OTDR to operate independently requires a new set of optical passives that have sufficient isolation and minimal insertion loss to enhance sensitivity across the board. This is described next.


## **7. Role of Optical Passives**

Architecting a pervasive monitoring infrastructure requires one to follow the fundamental rules laid out in earlier sections. It requires not only the ability to accommodate existing, set optical passives and access architectures, but also to define reference optical passives and architectures that can be better served. To this end, we have specified new sets of optical passives with consolidated test points.

Typical optical passives have had multiple filtered ports that either take in or put out ITU WLs. Then they have a COM port that combines these to launch on the optical fiber. WLs within the ITU range but that are not used in the Mux are typically available in the Upgrade port (UPG), whereas wavelengths outside of the ITU band are available in the Express (EXP) port. Typically, a test point for forward wavelengths is available (TPF) as well as a test point for reverse WLs (TPR).

Use of optical passives of this kind typically is less than optimal for pervasive monitoring. To peer in the network at 1610 nm, one would have to use the whole of the EXP; to check on outgoing WLs, one would have to use all of the TPF, and for the return WLs, the whole of the TPR. Doing this would require 3 ports on the monitoring tool, but more importantly it eliminates the future ability to add out-of-band wavelengths to the link, and requires manual troubleshooting, should the need arise, by the using TPF or TPR, since they are all connected to the tool. While this already considerable base of passives can also be monitored subject to the limits discussed, we opted to define a new set of optical passives to improve on this efficiency in a significant way.

Furthermore, optical switches are typically used for one WL at a time, and typically to direct light from one end to the other. Since all but one port is active, interactions between ports that are not in the optical path are generally not taken into account to build inexpensive switches. In our case, however, ALL ports that are lit on the switch are lit, which can cause unintended interaction between unrelated ports that are not in the intended light path. This is a fairly complex subject and not discussed further, but fortunately, optical passives can be designed quite easily to avoid this crosstalk issue altogether.



**Figure 7 – Optical Passives Old and New in Comcast**

The new standard of optical passives in Comcast today have consolidated test ports so that manual troubleshooting, if needed, can occur independently of pervasive monitoring. This is accomplished through an innovative optical filter design that consolidates the forward, reverse and 1610 nm part of the optical spectrum in one test port called TPFRO. The manual troubleshooting port is a consolidated forward and return test port called TPFR. The remaining optical wavelengths in EXP and UPG are all together combined into one consolidated Auxiliary port (AUX), which extends from 1260nm to 1598 nm (excluding the water peak WLs). This patent-pending concept is shown in Figure 7.

To enhance density as we prepare for the access growth, we have moved towards very small form factor (VSFF) connectors. These connectors can enhance deployment density, which is always useful in cramped PHEs and SHEs, thereby making space for pervasive monitoring tools.

## 8. Infrastructure Evolution

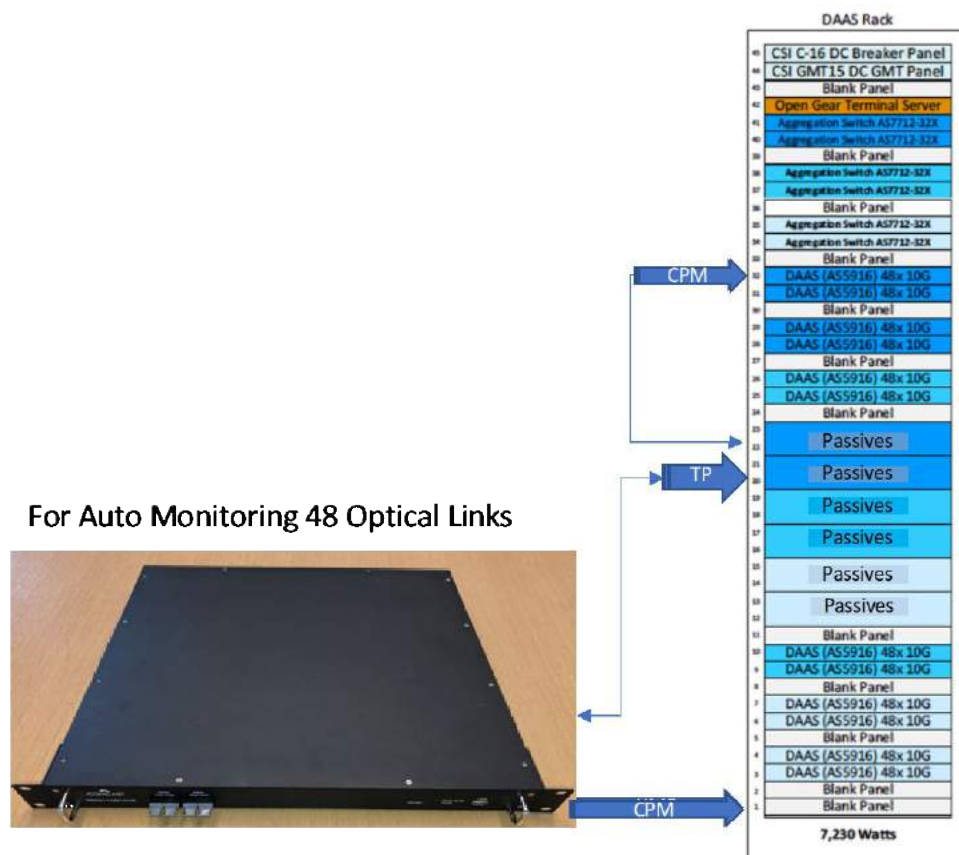
The Continuous Pervasive Monitoring tool (CPM) is implemented as a 1RU box that comprises the OSA, OTDR, 48-port optical switch and associated power supply and processors. As discussed, the processor must be fast enough to process fiber lengths, put out events, process the OSA and analyze wavelength information within the dwell time. The high reliability 48-port optical switch typically switches from port to port within 100ms. Therefore, within about 3 minutes (180 seconds), one can process up to 48 optical links, confirm their integrity as well as process up to 2034 wavelengths and confirm their well-being. It's a good time to be a fiber wavelength.



**Figure 8 – Actual Comcast Installation of the Continuous Pervasive Monitor**

As shown in Figure 8, the CPM is connected to the headend optical passive test points, which are then connected to DAAS ports. Figure 9 shows a typical SHE layout that accommodates the CPM.





**Figure 9 – Describing the DAAS Pod and the Monitor**

As indicated in Figure 9, the (Continuous Pervasive Monitor) CPM is typically set up in the SHE alongside the DAAS pods. These pods were well defined and have their own power supply panels, aggregating switches if needed, and DAAS switches, along with optical passives, in one compact location. The DAAS pods connect to the vCMTSs that are typically located in the PHEs. From 192 to 576 service groups (SGs) may be served with one DAAS Pod. The CPM is thus exceptionally well suited to monitor each of these DAAS pods and additionally handle any additional service that may be typically lit up in the SHE.

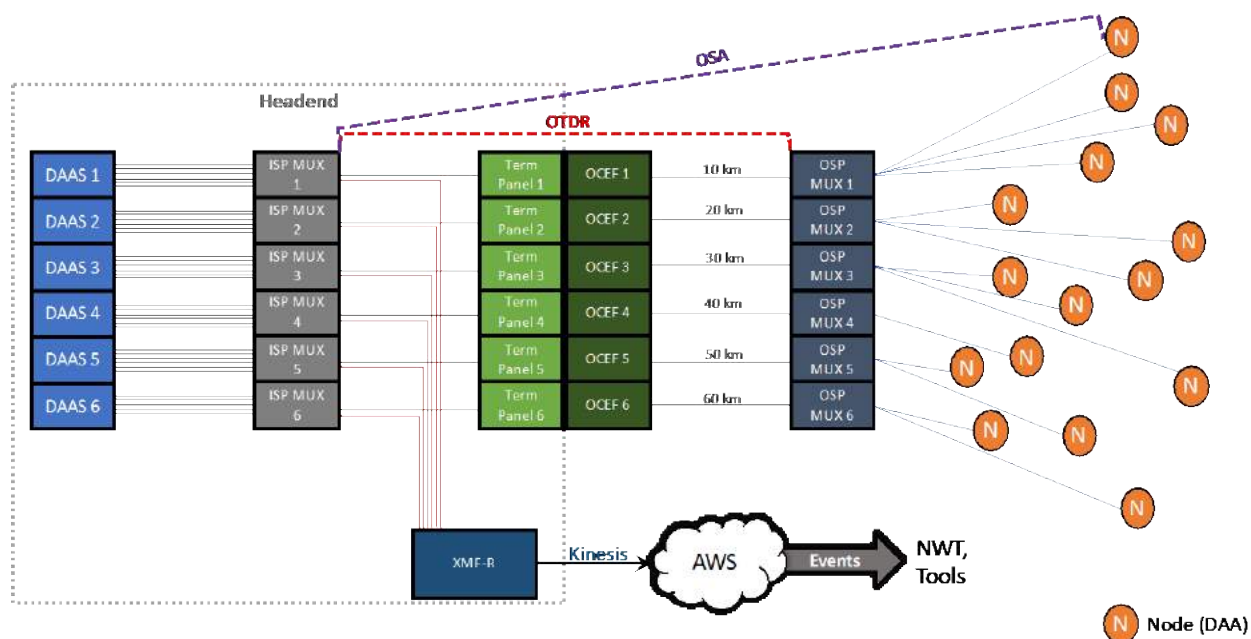
Typical SHEs may contain DAAS pods, many MetroE service groups, some traditional analog links and the incoming high capacity coherent links. Some SHEs are co-located with the PHE (they are called direct feed links), while many are small, single room discrete locations within the footprint. Some of the SHEs could be smaller and be accommodated in large air conditioned cabinets.

Widespread use of CPMs in the system will also help with SLAs in MetroE links and enable us to enhance the customer experience for both residential and business services. As such, even though the CPMs are never in the signal path, they still undergo a stringent Physical and Environmental (PnE) evaluation process, to assess their suitability of existing in SHEs and

meeting the demanding standards of our field equipment. Furthermore, as will be seen in the following section, and since they are going to be connected to our network, they undergo a rigorous security audit to ensure that there are no open ports, and that their connections are secure. It was surprising to see how quickly open ports are pinged by untrustworthy IP addresses, within hours of installation, making one feel glad about the security process in place.

## 9. End-to-End Optical Architecture

We have so far described the process of building innovative and cost-effective monitoring equipment, and some of the infrastructure, such as optical passives, that enhance its usefulness. Creating a physical infrastructure is one thing, but creating an end-to-end software infrastructure based on the third insight -- to take the promise of pervasive monitoring -- and turning it to reality is another substantial endeavor, and demonstrates the great opportunity to bring complex ideas into existence.



**Figure 10 – End-to-End Software Paradigm**

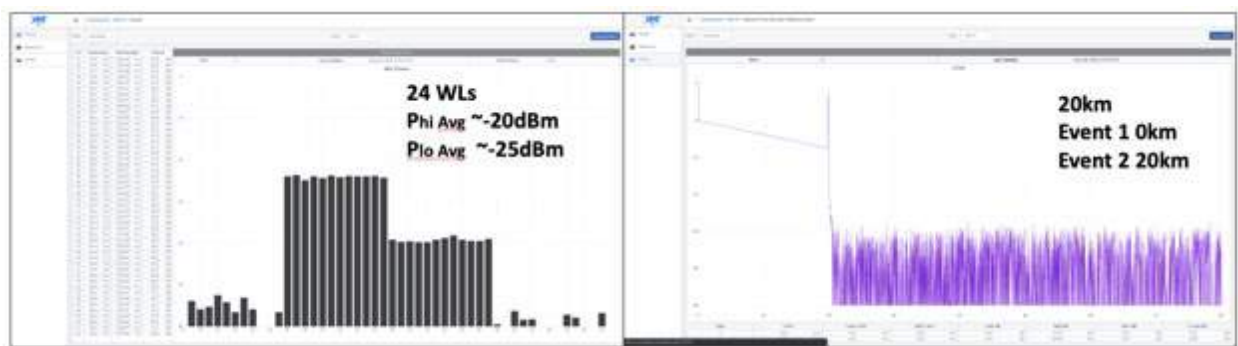
Figure 10 shows an end-to-end architecture in the context of our DAA architecture. As explained earlier, in the SHE, DAAS ports are connected to respective optical passives. Each DAAS port is logically connected to one optical node in the field and comprises two wavelengths, one downstream and one upstream. We also now use tunable SFPs in DAA applications that are typically spread from ITU 14 thru ITU 61. The light from multiple DAAS ports is multiplexed at the optical passives and is then connected to a termination panel in the SHE. This termination panel is the demarcation point between the SHE and the outside plant. Each fiber that goes to the outside plant has one optical connector on the termination panel. The termination panel fiber then passes thru a large conduit, called the Optical Cable Entrance Facility (OCEF) to the outside plant. Once outside, fibers are distributed according to plan to the various nodes. In our DAA

architecture, we use only one single fiber for US and DS operation. This makes the process of keeping track of fiber easier. With modifications, architectures that envision separate US and DS fibers can also be similarly monitored.

Typical Comcast links could be as long as 60 km, although the vast majority of our links for this type of architecture are around 30 km. After traversing the distance from SHE to node thru several splice enclosures, the fiber reaches the outside plant (OSP) mux. The fiber route is contained within our GIS systems, in special software entities such as SNET or Bentley. As the fiber makes it way to the OSP Mux, we recognize that its weight causes the fiber to sag between poles in aerial plant. And, every 1,000 feet or so, about 150ft of slack cable is maintained to aid in fiber cut restorations. Finally, as the fiber is lashed, the lashing wire envelops the strength member in a helix pattern. Combined together, these three factors contribute almost 20% extra linear fiber, as compared to the ground distance on maps. Many GIS systems take this into account, however there is an opportunity to re-verify and fine tune these numbers using the CPM devices.

In the OSP Mux enclosure, the individual wavelengths are brought out and fibers connected to support each node. Since the OTDR is a 1610 nm OTDR, once we reach the Mux, and the individual wavelengths are separated, that also ends the OTDR's ability to peer any further into the network. Therefore, the OTDR system provides continuous monitoring from the Inside Plant (ISP) Mux to the OSP Mux. Farther from the Mux to each individual node are set two fibers, one for US and one for DS, and these are terminated in the node to the SFP. This SFP receives the optical input from the DAAS port SFP and sends out light to be received by the DAAS SFP, thus completing the circuit. Notice, however, that unlike the OTDR, the ability of the OSA to view DS and US wavelengths is unaffected. It is able to record all wavelengths that traverse the fiber. Thus, a combination of OSA and OTDR provide a continuous and end-to-end view of the entire DAA network.

In keeping with our objectives, the CPM is connected to the TPFRO test point of the ISP Mux. The CPM is thus in a great shape to dwell on each one of its ports, measure the input and output wavelength values and their power values, independently measure the fiber link, and record any events on the fiber link. Using the fast processor with the unit, any event information can be uploaded to the cloud.



**Figure 11 – Sample view of the OSA and OTDR Output for each port**

## 10. Praemonitius est Praemunitious (Forewarned is Forearmed)

While information is uploaded to the cloud continuously, it is also used to train the monitor. This step is described later in the paper. Any changes to the optical levels or to fiber events are then reported via the automatic feed (called Kinesis) to an eventing engine that connects to a sophisticated internal resource called the National Watchtower (NWT).

Since outage information from various sources arrive at the decision engine with various minor delays, if the decision engine begins reacting immediately, it will have reacted with insufficient information. Thus, to avoid this “race” condition, the decision engine “soaks” or stores all incoming information with various minor delays to look at it holistically [4]. Then they are brought up as one unified ticket with actionable data and sent to respective locations. Ideally, if the contents indicate a fiber cut, these are also laid out on the GIS, and a street address location is also provided, in which case the repair crews also have a specific location to head out to. For this reason, slack/sag/helix and other such parameters are estimated from the appropriate databases to ensure a level of accuracy.

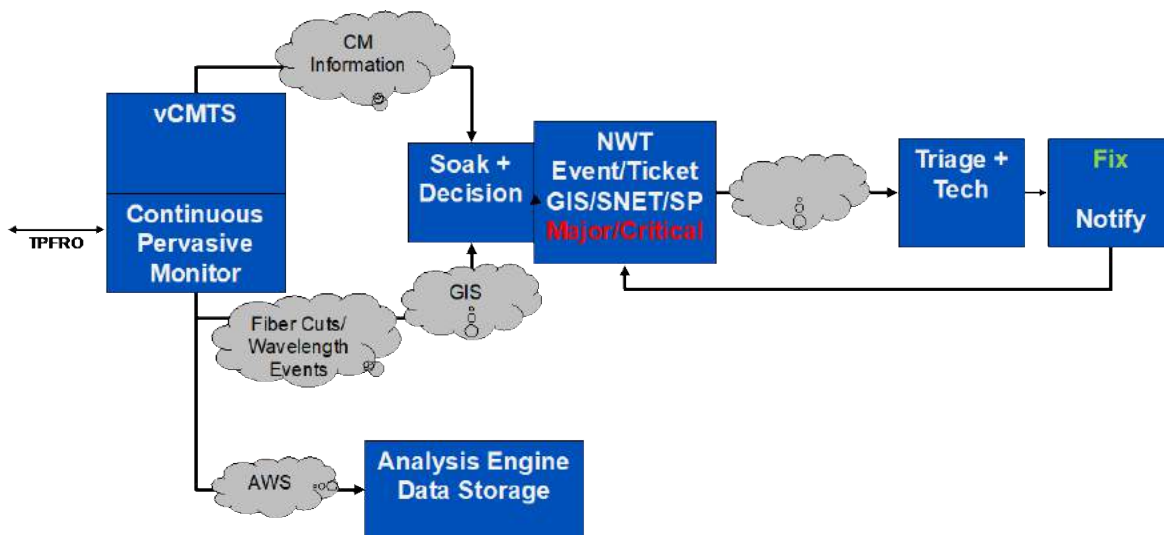


Figure 12 – Sample view of the OSA and OTDR Output for each port

We began this paper by noting the energizing effect of a suspected fiber cut. This process not only establishes the said cut, but also provides a specific location in almost real time. In Figure 12, it can be seen that the events are sent to the decision engine for handling related to ticketing from the CPM. To aid in fiber cuts, we also estimate the distance of the fiber cut and its relative location based on sag/slack/helix and compared to the GIS. At the decision location, this information is soaked and compared to other information also flowing, to avoid race conditions. At the NWT the ticket is generated with an appropriate alarm level, and the location of the fiber cut/node location is then sent to triage and to the technician. The problem is then fixed, and the tech clears the issue.

While events are sent to a decision engine, the periodic data collected is sent to the cloud continuously. This data is curated and available for deeper analysis. For example, this data may be examined to compare various links at locations at the same time, or to compare data from the same link during different parts of the years, for weather-related impacts. Over time, this data will be a treasure trove for a machine learner to sift thru for various patterns in optical preventative network management.

If, on the other hand, wavelengths disappear on the OSA, but the fiber link on the OTDR shows strong, the inference is that the node is individually affected. This could mean that the node has an electrical, RF or optical issue inside of it, or that the fiber between the node and the OSP Mux is cut. In this case, after a soak period, the affected node is identified, and the technician can easily figure out the node status. If the node is fine, then the fiber cut between the OSP Mux and the node is easily identified by an OTDR (this will require a handheld OTDR, which is a mobile equivalent of the rackmount CPM). One other option, which is a bit manual, requires disconnecting the headend DAAS port of the affected node and shooting the fiber with a tunable OTDR, but this process requires a fairly good view of the ISP optical connections.

Establishing and then maintaining end-to-end connectivity for all elements in the network is as complex as it is important. The OSP plant that begins from the OCEF is typically overlaid on the GIS, as described earlier. This can span several 10s of km of plant. Also, such connectivity must be maintained within the PHE or the SHE, which is considerably difficult. This is because equipment may be periodically rearranged, revamped and swapped out, sometimes in phases. Therefore, an auto discovery process for all field entities is the most optimal path, which begins with a robust provisioning process.

Because the data is stored in an efficient format in the cloud, this data can be compared year to year to understand how each of the fiber links perform in different weather conditions. The role of machine learning in this type of an effort cannot be overestimated. Through a systematic comparison of events and fiber profiles across the fiber lengths, not only can link performance be flagged, but weaker splice enclosures prone to fiber fills can be replaced to shore up the fiber availability. In a natural disaster, having access to multiple fiber cuts and their location in real time will enable a far more efficient dispatching of resources. Since optical passive information is stored along with the current usage of optical wavelengths, and available at any given time, capacity on each fiber link can be estimated in real time. This helps us to understand which links can handle additional wavelengths. Furthermore, some of our limits on optical level differentials and wavelength plans to accommodate analog, coherent and 10G wavelengths on the same fiber can not only be viewed but also enforced.

The CPM is deployed in all three divisions at Comcast. The cloud-based infrastructure component, with respect to its provisioning and continued use, is described next.

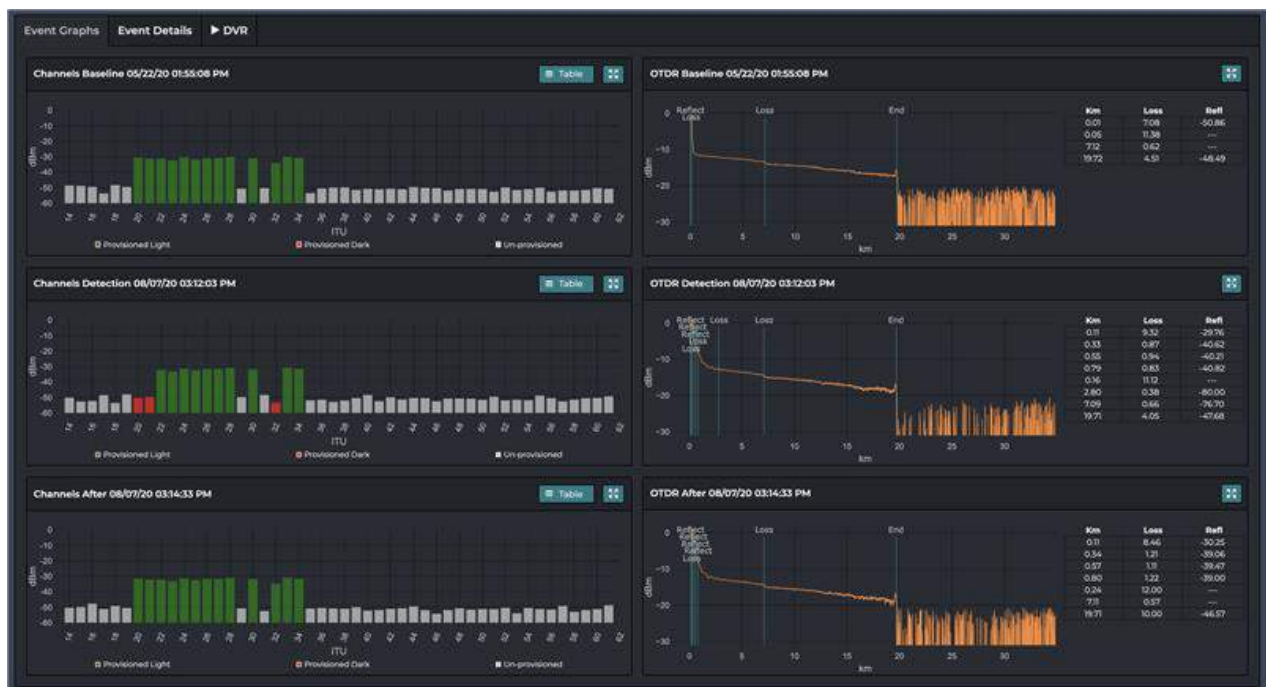
## **11. Provisioning and Training the Monitor**

When the CPM is first deployed, the MAC is discovered, an IP address given, and a name established. In the provisioning process, the ports are assigned to each of the ISP Muxes, by

[illegible]

Once this is established, then the CPM is left alone until it has acquired, by the port dwelling process described earlier, a baseline of its fiber length and the wavelengths. Once the baseline is verified to be correct according to the maps, the baseline values are set, and then subsequent values are compared to the baseline in perpetuity. This is indicated by the green color of the OSA wavelength representation. In a typical day, each link is verified 480 times (or more per 24 hours), at three minute intervals. Thus, any change in the current vs. baseline situation is immediately flagged and sent to the NWT for ticketing.





**Figure 14 – Fully Provisioned Port (notice the green and red WL colors)**

Figure 14 shows a case where the baseline has been set, as can be seen with all of the green ITU wavelengths. In the middle, we see that a power outage has taken out three nodes. Because in this case the CPM is monitoring the US ITU wavelengths, we see three of the green wavelengths are now in red. However, a quick check at the OTDR confirms that the trunk fiber is not cut -- it shows no events. We would be able to verify thru Continuity (an internal “source of truth” dashboard that indicates power continuity in the network) about the power outage. Once the power is restored, the wavelengths come back on-line and compare favorably to the baseline. In this case, this outage was resolved without the need to visit the node. However, if the power had been on, then the next step would have been to understand if there was a fiber cut between the OSP Mux and the nodes (although this is unlikely, given that more than one node was affected). In that case, a technician would resolve the issue by means of more sophisticated tools, described in the next section.

Occasionally, it becomes necessary to interrupt the automatic process to troubleshoot or verify performance. In these cases, the software also lets one SSH into the CPM to direct it to specific ports, out of the order for immediate real time view. In these cases, the CPM understands the legitimate request and pauses the auto view.

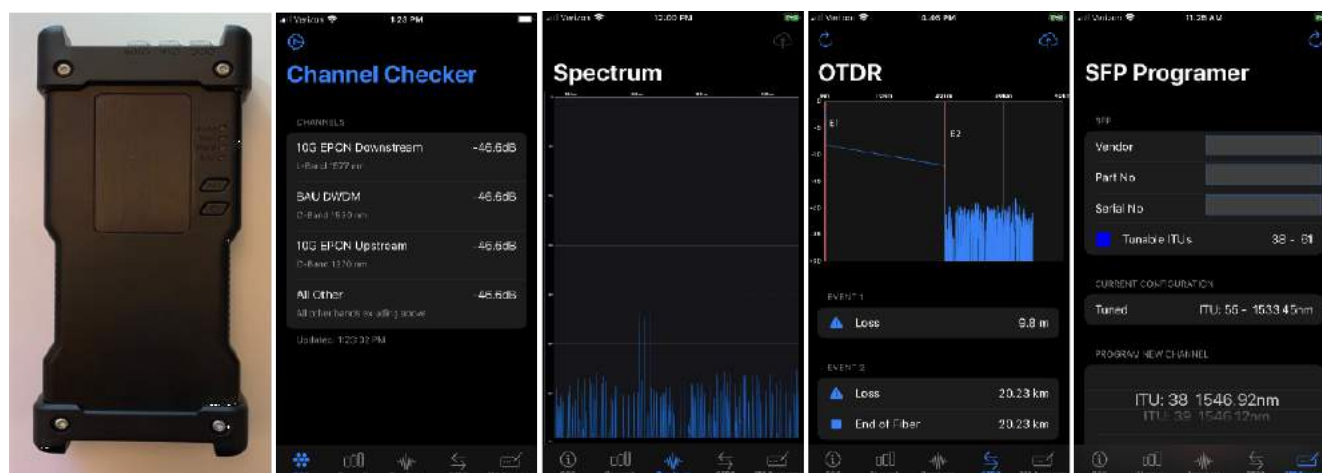
This level of automation and remote capability was especially helpful in the COVID environment, when folks from multiple locations were suddenly able to get on a conference call to view, inspect and troubleshoot specific links. In one specific instance, we were perplexed by a rather high loss appearing in the dead zone of an OTDR. The CPM was provisioned, by alerts that were not yet firing. Perplexed by this, and by good, old-fashioned process of elimination, we figured out that because of the unusually short fiber length, and to protect the APD of the SFPs, a rather large attenuation pad had been placed in the COM port of the ISP passive -- thus

contributing to the dramatic attenuation increase on the OTDR and the OSA. Thus, it is rather important to understand the role of attenuators (commonly used to attenuate light levels) and their effect on OTDR resolutions and OSA power levels.

## 12. Handheld to the Cloud

Thus far, we have described the role of continuous and pervasive monitoring of optical assets. However, many impairments in the system require human interaction and experience in order to reach a resolution. In these cases, we rely on the considerable acumen and insights of our headend and fiber technicians. We also outfit them with an array of sophisticated optical equipment. These could be full featured optical spectrum analyzers, high sensitivity and/or high input optical power meters, 1550 nm tunable OTDRs and SFP programmers. We have previously mentioned that the optical passives have been designed to enable manual as well as pervasive monitoring.

In the example in the previous section, if only one node was affected, the technician would have disconnected the port to the affected node and used a tunable OTDR to verify whether there was a fiber cut, as it can peer beyond the Mux. Alternatively, a technician would have been dispatched to the node location for troubleshooting.



**Figure 15 – Handheld unit with a Mobile Application and Cloud connectivity**

The previously mentioned handheld equivalent of the CPM is also potentially useful for field technicians. This device has a C-Band OSA, the 1611 nm OTDR, an SFP Programmer and a Channel Checker/Power meter to aid in PON deployments. Designed with a long battery charge (a full shift) and WiFi hotspot, this device enables field technicians to troubleshoot a variety of fiber cuts, and SFP installations, along with routine troubleshooting. Based on popular RF meters of the same type, this device has impressive cloud capabilities, with a UI designed in-house. Because all of our technicians carry iPhones, the UI and display is adapted to the iPhone and shown in Figure 15. A device of this kind allows the technicians to troubleshoot and upload their observations to the cloud and helps round out our optical asset monitoring.



## 13. Conclusions

We began this paper with a description of the access footprint and our substantial efforts to address continuous and pervasive monitoring of our optical networks. An innovative mix of technology, in use in other areas of optical communications, including innovations in optical passives was critical in standing this system up. As important as the hardware is, the software infrastructure is even more so.

Being able to remotely and automatically get notifications of fiber cuts, with appropriate overlays on GIS maps, is extremely useful in reducing the time to repair. It is especially helpful to not only DAA and HFC links, but also to the MetroE links in the access domain. The handheld unit offers an extra layer of efficiency across the board and helps fulfill the promise of continuous and pervasive monitoring of access optical assets. The benefit of CPM during the COVID pandemic has been rather positive and can potentially be generalized further to enable remote monitoring and mitigating capabilities across the network.

## 14. Acknowledgements

It is a pleasant duty to acknowledge the entire team within Comcast that has been working directly and indirectly on the Continuous Monitoring project. A project of this magnitude benefits from the dedication of diverse expertise from hardware and software at the Operations and technology team, the reliability and functionality testing of the Physical and Environmental test team, the ticketing and alerting of the NWT team and the functional rules for deployment from the access engineering team. Our thanks are especially due to our vendor and partner II-VI for their insights. We sincerely thank the Senior Leadership Team at Comcast NGAN in supporting this project and for their support in deploying it in all the three divisions at Comcast.

## Abbreviations

|      |                                        |
|------|----------------------------------------|
| 4WM  | Four Wave Mixing                       |
| APD  | Avalanche Photodiode                   |
| BDR  | Baseband digital receiver              |
| CMTS | Cable modem termination system         |
| CPM  | Continuous Pervasive Monitor           |
| CPU  | Central Processing Unit                |
| DAA  | Distributed Access Architecture        |
| DAAS | Distributed Access Architecture Switch |
| DS   | Downstream                             |
| DWDM | Dense Wave Division Multiplexing       |
| EXP  | Express port                           |
| FPGA | Field Programmable Gate Array          |
| FS   | Full Spectrum                          |
| GIS  | Global Information System              |
| GOA  | Grey Optics Aggregation                |
| HFC  | Hybrid Fiber Coax                      |

|       |                                        |
|-------|----------------------------------------|
| HHP   | Households passed                      |
| ISP   | Internet Service Provider              |
| ITU   | International Telecommunications Union |
| MAC   | Media Access Control                   |
| NWT   | National Watch Tower                   |
| OCEF  | Optical Cable Entrance Facility        |
| OEM   | Original Equipment Manufacturer        |
| OSA   | Optical Spectrum Analyzer              |
| OSP   | Outside Plant                          |
| OTDR  | Optical Time Domain Reflectometer      |
| PHE   | Primary Headend                        |
| PON   | Passive Optical Network                |
| RF    | Radio Frequency                        |
| ROADM | Remote Add-Drop Multiplexor            |
| RPD   | Remote PHY Device                      |
| SFP   | Small Form Factor Pluggable            |
| SG    | Service Group                          |
| SHE   | Secondary Headend                      |
| SLA   | Service Level Agreement                |
| SOAP  | Switch On A Pole                       |
| SSH   | Secure Shell                           |
| TPF   | Test Point Forward                     |
| TPFR  | Test Point Forward and Return          |
| TPFRO | Test Point Forward, Return and OTDR    |
| UPG   | Upgrade Port                           |
| US    | Upstream                               |
| VSFF  | Very Small Form Factor                 |
| WL    | Wavelength                             |

## Bibliography & References

1. Comcast's Extensive Nationwide Network: <https://business.comcast.com/about-us-backup/our-network-backup#:~:text=Comcast%20Business%20Fiber%20Optic%20Network,than%20145%2C000%20miles%20of%20fiber.>
2. Comcast Access Footprint: <https://www.cabletv.com/xfinity/availability-map>
3. *Accelerating the Virtualization: Introducing the Hybrid Fiber Shelf into the Mix*, Venk Mutalik, Bob Gaydos, Dan Rice and Jorge Salinger, SCTE EXPO 2020
4. *Operationalizing the Grey Optics Architecture: An Update a Year After*, Venk Mutalik, Dan Rice, Bob Gaydos, Doug Combs and Pat Wike, SCTE EXPO 2020

# **Connectivity and COVID-19: Maintaining QoE During a Crisis**

Technical Paper prepared for SCTE•ISBE by

**William McFarland**  
CTO  
Plume  
290 California Ave., Palo Alto, CA 94306  
650-823-6315  
Bill@Plume.com

# Table of Contents

| <b>Title</b>                                                               | <b>Page Number</b> |
|----------------------------------------------------------------------------|--------------------|
| 1. Introduction .....                                                      | 4                  |
| 2. The Effect of COVID-19 on User Behavior.....                            | 6                  |
| 2.1. Changes in Working at Home .....                                      | 7                  |
| 2.2. Changes in Usage .....                                                | 9                  |
| 3. The Effect of COVID-19 on Networks and Technologies to Compensate ..... | 10                 |
| 3.1. Demands on Coverage, and Multi-AP Solutions.....                      | 10                 |
| 3.2. Increased Load, and Throughput Optimized Steering.....                | 13                 |
| 3.3. Interference and MDU Joint Optimization.....                          | 14                 |
| 3.4. Cyber Security Attacks and IoT Device Security.....                   | 17                 |
| 3.5. Traditional QoS vs. QoE .....                                         | 20                 |
| 4. The Effect of COVID-19 on Financials and How to Compensate.....         | 22                 |
| 4.1. Proactive Support.....                                                | 23                 |
| 4.2. Lost Revenue and Compensation with Additional Services .....          | 24                 |
| 5. Conclusion .....                                                        | 25                 |
| Abbreviations.....                                                         | 25                 |

## List of Figures

| <b>Title</b>                                                                      | <b>Page Number</b> |
|-----------------------------------------------------------------------------------|--------------------|
| Figure 1 - Speed of Broadband Access to Homes Before and After COVID-19.....      | 5                  |
| Figure 2 - Load and Usage Before and After COVID-19.....                          | 6                  |
| Figure 3 - Increase in Working From Home in the US.....                           | 7                  |
| Figure 4 - Working From Home in Canada.....                                       | 8                  |
| Figure 5 - Working From Home in the EU .....                                      | 8                  |
| Figure 6 - Hours of Active Time by Device Type.....                               | 9                  |
| Figure 7 - Coverage Alarm vs. GW only or GW + additional APs.....                 | 11                 |
| Figure 8 - Single AP vs Traditional Mesh vs Optimized Adaptive Wi-Fi.....         | 12                 |
| Figure 9 - Optimization System for Wi-Fi Networks .....                           | 13                 |
| Figure 10 - Throughput Optimized Steering.....                                    | 14                 |
| Figure 11 - Interference Histogram for Suburban Homes and Urban MDUs.....         | 15                 |
| Figure 12 - Before and After MDU Joint Optimization Interference Histograms ..... | 16                 |


|                                                                             |    |
|-----------------------------------------------------------------------------|----|
| Figure 13 - Cyber Security Attacks Before and After COVID-19 .....          | 17 |
| Figure 14 - Scatterplot of Tx and Rx Bytes per Minute for Nest Camera ..... | 18 |
| Figure 15 - Lateral Movement of Viruses and Security at Every Node .....    | 19 |
| Figure 16 - QoE Factors and Outputs .....                                   | 21 |
| Figure 17 - Struggling Devices Identified by QoS and QoE.....               | 22 |
| Figure 18 - Support Call Prediction Precision vs. Recall .....              | 23 |
| Figure 19 - Wi-Fi Motion Detection Concept.....                             | 24 |

# 1. Introduction

Few events in modern times have had as large an impact on society as COVID-19. Across the globe, people have been encouraged to work- and school-from-home. In many cases, this was mandated by governments by shelter-in-place orders –and for some, this continues to be the ‘new norm’. Not surprisingly, this has resulted in a huge shift in network usage patterns. Several trends are apparent in the data that will be presented in this paper:


- Large increases in people working on the internet at home during the workdays
- Dramatic increases in the number of hours that devices are active on the home network, particularly computers, phones, and entertainment devices
- The workday is spreading into the evening hours as families try to juggle work and family commitments

Remarkably, wired internet access networks have held up to the added load well. **Figure 1** shows the download and upload speeds as recorded at more than 15 million households that are managed by the Plume Cloud. The delivered speeds, as measured by Plume, are remarkably consistent before and after COVID-19 caused the large increase in working-from-home. In fact, average speeds even trended *upwards* across the time due to subscribers upgrading to higher tier services to meet their expanded needs. At the onset of the pandemic many predicted that broadband service providers would struggle to meet higher and changing bandwidth demands—many OTT TV providers *even* reduced streaming quality to assist—however, Plume data shows that the networks coped admirably,



**Figure 1 - Speed of Broadband Access to Homes Before and After COVID-19**

While service providers have generally been able to keep up with the added load in delivering service to the edge of the customer’s home, the story within the home could be much different. The majority of devices in homes today connect over Wi-Fi. Wi-Fi always has the fundamental problem of being a shared, and oftentimes best-effort medium, so handling increased load can be particularly challenging. **Figure 2** shows the dramatic increase in load from three key categories of devices: phones, computers, and entertainment devices (set top boxes, TVs, and gaming consoles). The increases are shown for both the amount of data consumed by these devices, as well as the number of minutes they were active on the network. In the stacked graph, the values are normalized to highlight the percent change from before to after COVID-19.



**Figure 2 - Load and Usage Before and After COVID-19**

The key takeaway from **figure 2** is that inside the home, a high degree of “wireless chaos” has been triggered by COVID-19. The remainder of this paper looks at this phenomenon in more detail, and describes what service providers can do to overcome this chaos and provide their customers with a highquality of *experience*.

## 2. The Effect of COVID-19 on User Behavior


As pointed out in the introduction, COVID-19 has changed network usage patterns significantly. This section examines these changes in more detail. While it is impossible to predict for sure, it is likely that



some of this behavioral change will persist long after COVID-19 has receded. If it recedes at all. For example, numerous companies, in the wake of the surge in working at home with COVID-19, have announced permanent acceptance of working from home. It is likely that many employees will take advantage of these offers. While this paper focuses on the changes before and after the onset of COVID-19, it is likely that the new behaviors, and the required adaptations by service providers for them, will be long term.

## 2.1. Changes in Working at Home

The most immediate effect of COVID-19 is an increase in the number of people working from home. Plume was able to analyze the data collected by Wi-Fi networks that it manages for service providers around the world to identify the change in the number of people who are working-from-home. The data analyzed included the type of devices, amount of data, amount of active time on the network, and domains accessed. Combining these observations, Plume is able to gauge homes in which at least one person is working from home during the 9am to 5pm weekday period. **Figure 3** shows the result for fourteen major metropolitan areas in the US:




**Figure 3 - Increase in Working From Home in the US**


While the reaction across cities is relatively consistent, some details are instructive. The rise in working from home (WFH) occurred first in Seattle, the first city to see a significant number of cases, and the first

city to declare a shutdown. The San Francisco (SF) Bay Area had among the highest levels of extended WFH, perhaps enabled by its high percentage of tech workers. Cities in certain regions of the country (e.g. San Antonio and Jacksonville) have had somewhat lower levels of WFH. And we see some aborted attempts at loosening, followed by returns to higher levels of WFH.

Plume manages a large number of Wi-Fi networks in Canada and Europe as well as the USA. While the trends are similar, we can see a stronger, steadier downward trend in working-from-home in several Canadian cities and European countries (**Figures 4 and 5**). These correlate with countries that have been better able to control their COVID-19 case counts.



**Figure 4 - Working From Home in Canada**



**Figure 5 - Working From Home in the EU**

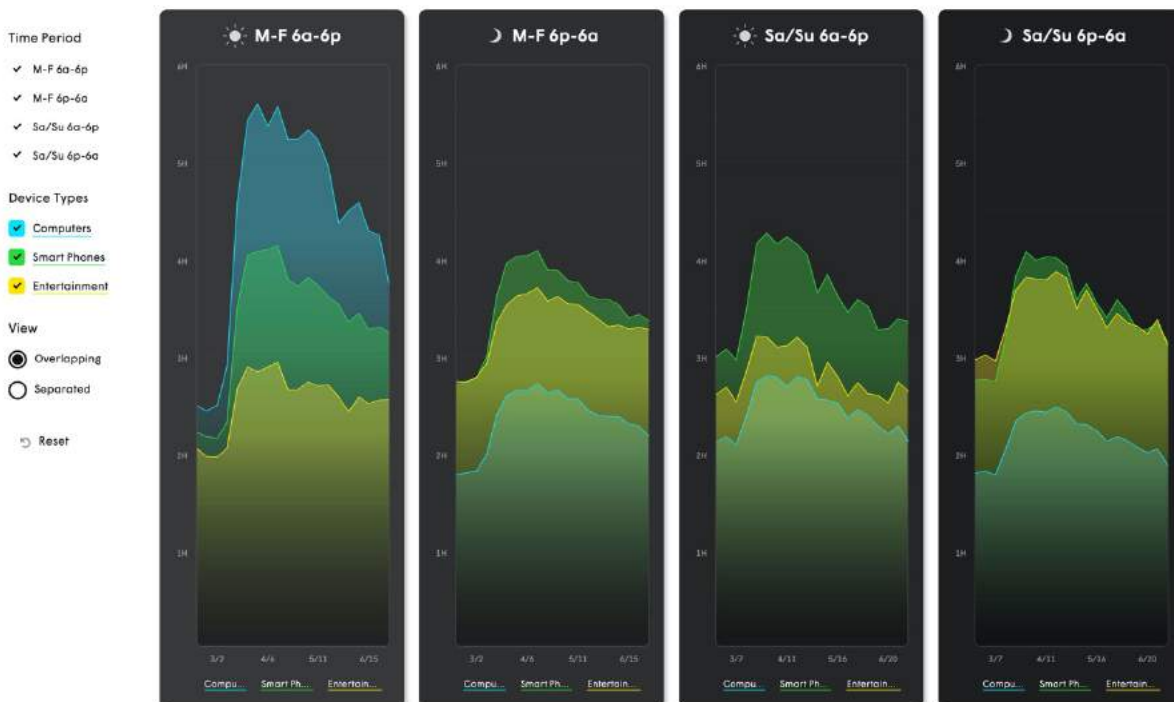
## 2.2. Changes in Usage

Increased WFH is the first link in the chain between COVID-19 and increased in-home network distress. The change in usage of devices is the second. Devices don't readily identify their device type when they connect to a Wi-Fi network. But Plume is able to use a variety of factors including DHCP requests, UPNP transactions, DNS requests, Host Names, and user agent (browser) transactions to identify device types. This is based on machine learning across tens of millions of homes that are operated by Plume. After identifying the device type, Plume is able to observe the amount of data consumed by the different device types, as well as the number of minutes that different device types are active on the network.

**Figure 6** shows Plume's observations regarding the number of active minutes by device type across homes in the US. The "Computers" category includes both desktop and laptop computers, and the "Entertainment" category includes set top boxes, TVs, and gaming consoles.

### Busy Hours at Home in the US

Entertainment Devices Include Set-Top Boxes, Smart TVs, And Game Consoles  
2/20 - 7/5 | DATA FROM 14 METRO AREAS, UPDATED WEEKLY



**Figure 6 - Hours of Active Time by Device Type**

By breaking the usage into days of the week and times of the day, several interesting before/after COVID-19 conclusions can be drawn. Unsurprisingly, the most dramatic increase in usage was for computers during the weekday work times. However, significant increases in phone and entertainment usage are seen in the evenings and on the weekends. This is likely due to the reduced entertainment options with social distancing. However, computer usage has grown significantly in the evenings and on weekends, indicating some spreading of working hours, as well as increased use of computers for entertainment.


### 3. The Effect of COVID-19 on Networks and Technologies to Compensate

The third link in the chain from cause to effect is how changes in usage due to COVID-19 impact the networks in people's homes. This section focuses on five of the most important effects. While these effects make the provision of adequate connectivity to provide high quality of experience more difficult, there are technologies that can compensate for the added challenges. These technologies are presented side by side with the description of the problem.

#### 3.1. Demands on Coverage, and Multi-AP Solutions

The most basic requirement for any Wi-Fi user is to be able to physically connect to the network. When people have trouble connecting in their own home, the problem is usually caused by inadequate coverage. Coverage is defined the extent to which a given home has adequate signal strength to and from all devices at every location in the home. Homes often have "dead spots," regions that are not covered well by Wi-Fi. Behavioral changes due to COVID-19 have exacerbated this problem. As multiple people in the home search out areas where they can work in privacy, they are more often trying to use regions of the house with poor coverage.

Since client devices typically transmit at lower power levels than the Access Points (APs), it is often the connection that comes back from the client device that is the limiting factor. Plume's networks observe the signal strength of transmissions coming from client devices, moving this information to the cloud. Based on analysis of this information in the cloud, it is possible to determine which homes have a coverage problem. For this analysis, Plume defines a home to be in coverage alarm if more than 25% of the client devices have a coverage issue, a coverage issue defined as a device with under -70dBm signal strength 50% of the time. -70dBm is the level at which data rates start to drop quickly, and clients start to spend a lot of time searching for better APs or roaming, making media flows less reliable. **Figure 7** below shows the percentage of homes that are in coverage alarm for a large North American service provider. As can be seen, more than 40% of homes that have only one Wi-Fi AP are in coverage alarm. However, by adding additional APs ("Pods"), that can be reduced close to 10%.

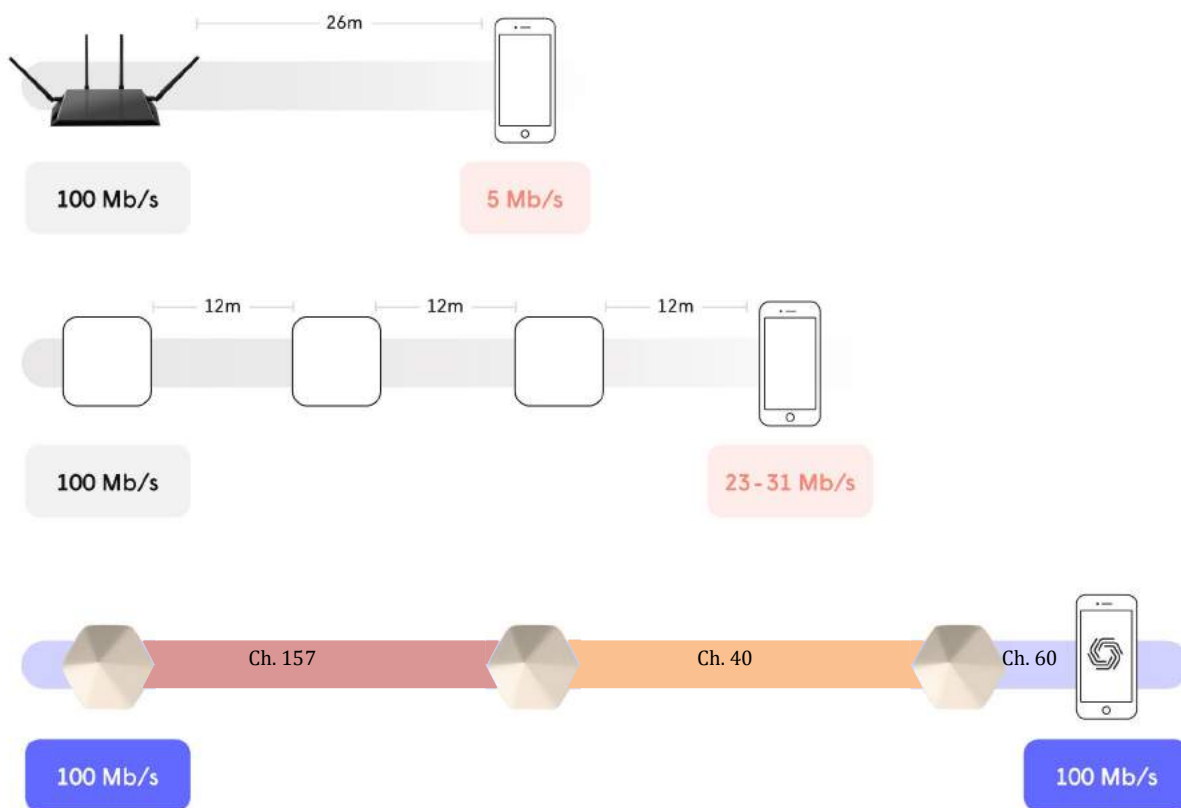


**Figure 7 - Coverage Alarm vs. GW only or GW + additional APs**

Typically, additional APs added into a home are connected wirelessly to each other. Few homes have Ethernet wires in locations that are appropriate for placing additional APs. Traditional Wi-Fi repeaters or mesh systems can solve the coverage problems just discussed, but they do so at the expense of throughput. These systems will utilize the same frequency channel for the backhaul connection between APs as is used to connect client devices to the AP (fronthaul). In fact, mesh systems often deploy the entire mesh, including all fronthaul and backhaul connections, on the same frequency channel. Such an arrangement suffers from self-interference in which transmissions on one hop in the network interfere with transmissions on other hops. For a two hop path, the throughput is more than halved, for a three or four hop path, the throughput is divided by more than 3 or 4 respectively.

This can be greatly improved if a different frequency channel can be used on each hop in the network.

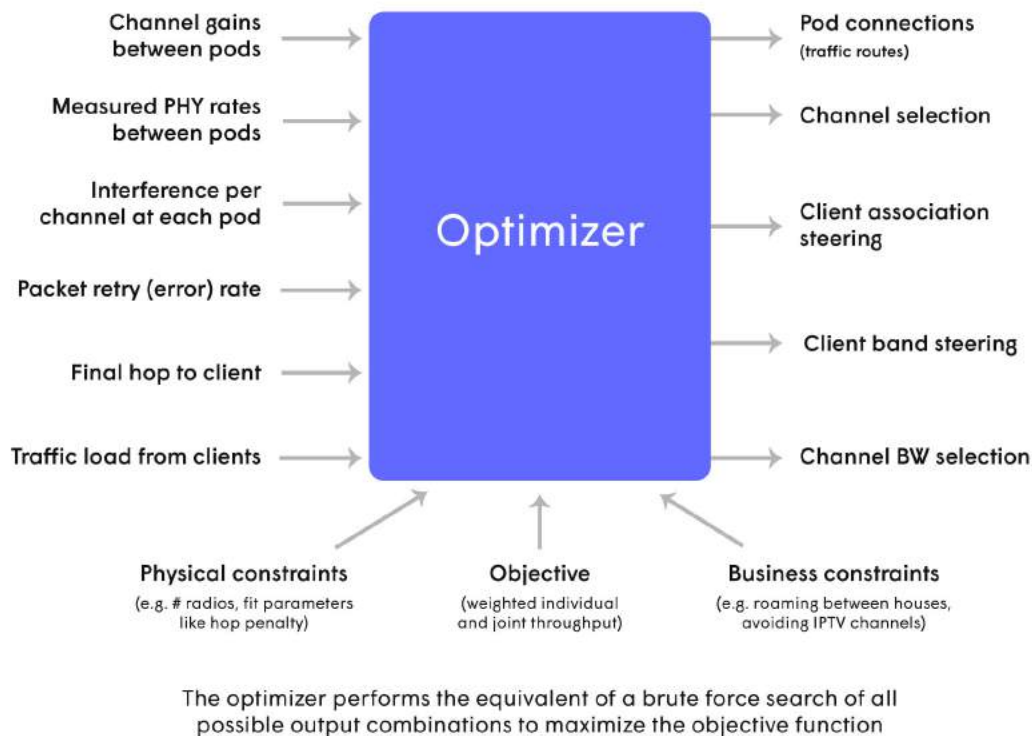
**Figure 8** shows a comparison between three approaches for serving a 20 meter connection in a typical home: direct connection to a powerful single AP, multiple hops through traditional repeaters or mesh system, and multiple hops through an optimized Adaptive Wi-Fi network. The difference in throughput achieved to the client is dramatic.



**Figure 8 - Single AP vs Traditional Mesh vs Optimized Adaptive Wi-Fi**

The complexity of choosing and configuring the optimized Adaptive Wi-Fi topology is why it is not often employed. However, modern cloud technology can enable sophisticated approaches. Rigorous optimization approaches, such as Mixed Integer Linear Programming (MILP), can be employed to maximize an objective function that includes throughput to individual clients, overall home system capacity, and fairness among the devices in the home. **Figure 9** shows a conceptual diagram of such an optimization system with its inputs and outputs. This approach is advantageous because the resulting topology deployed in a particular home will be optimal given the constraints of that particular home.






**Figure 9 - Optimization System for Wi-Fi Networks**

### 3.2. Increased Load, and Throughput Optimized Steering

The optimized Adaptive WIFI system presented in section 3.1 solves the problem of coverage while preserving optimal throughput. However, multi-AP systems in general bring the problem of where client devices connect into the network. Two problems are generally seen. First clients may be “sticky”, remaining attached to one AP even as they move to the far side of the home. This has led some manufacturers to implement “sticky client steering” in which clients that are detected to be “stuck” on far away APs are kicked off that AP, forcing them to search for a closer AP for connection to the network. While this is better than nothing, the connections achieved in this way are not optimum. And second, even clients that dutifully shift to the closest AP are not necessarily connected in the optimum way. This is fundamental, as the client cannot know the complete pathway through the Wi-Fi network to the Internet. All the client can know is the signal strength from it to the APs it might connect to. To achieve the optimum connection requires understanding the speed of the complete connection, perhaps through multiple hops, from the client to the Internet for each option. [In response to SCTE editor comment.]

**Figure 10** shows the desired outcome in client steering and AP topology. First, the optimizer eliminates excessive hops that would degrade performance. It does this for example by connecting APs directly to the GW AP as appropriate, when an extra hop to get through the network is not justified. Second, throughput optimized steering is used to move the client device to the AP where it will achieve the best

performance. The highest performance connection point for a client is often *not* the closest AP. As can be seen in the figure, connecting to the closest AP sometimes adds unnecessary additional hops, and can even increase the individual hop lengths in the connection.



**Figure 10 - Throughput Optimized Steering**

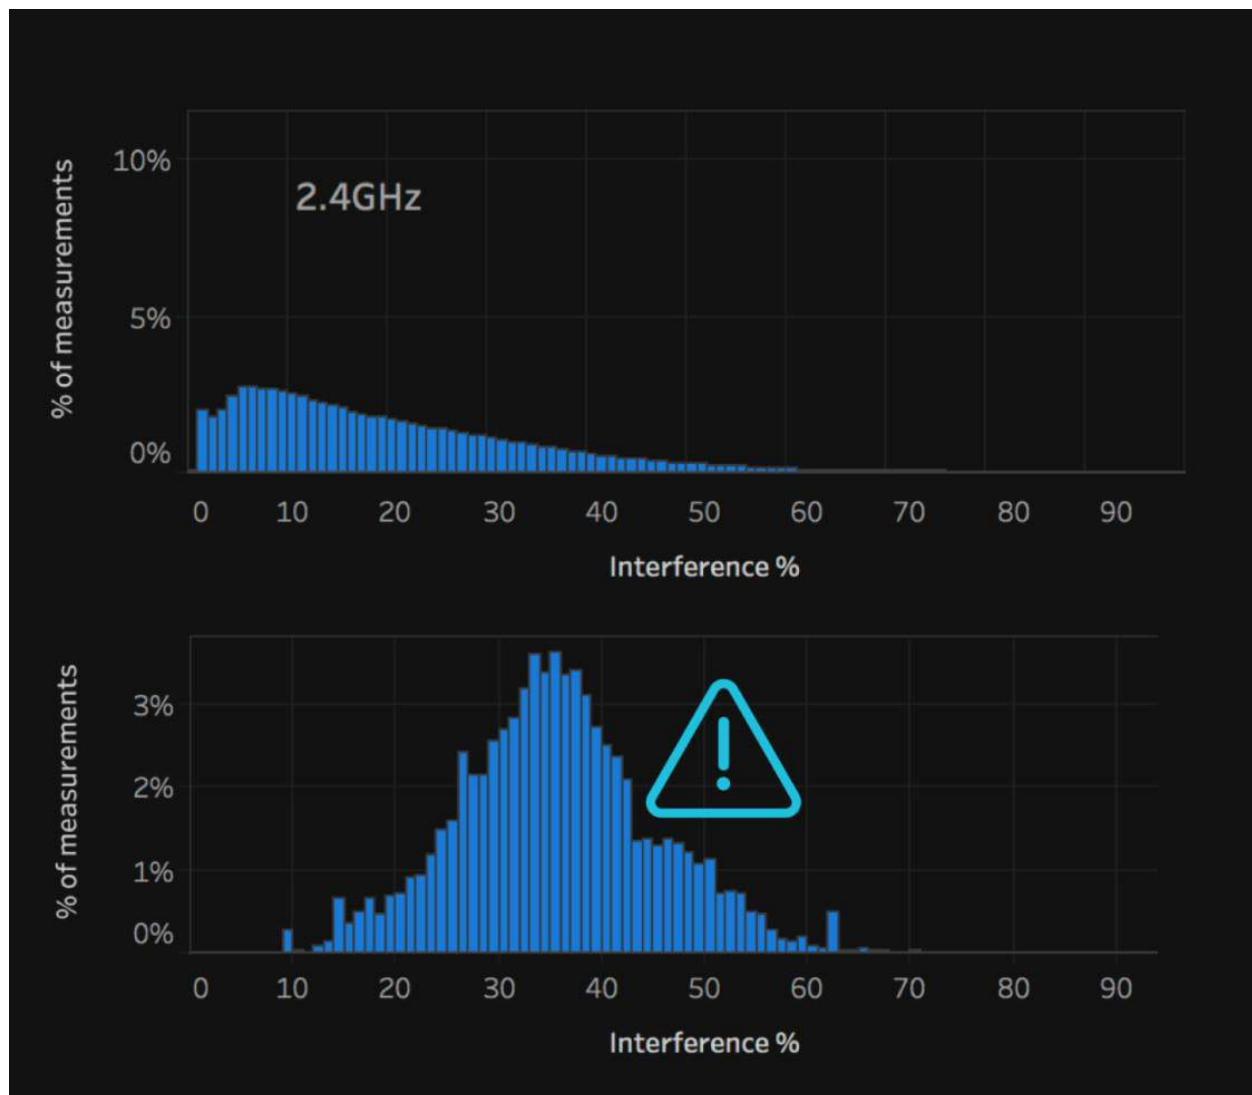
The best location for client devices can be calculated as part of the optimization procedure. As clients move about the home and attach in inappropriate places, the results from the optimization can be used to identify where they would perform best, and client steering can be used to move them to the correct location.

### 3.3. Interference and MDU Joint Optimization

Increased usage due to COVID-19 also causes an increase in interference. Both interference (overlapping transmissions from neighbors) and congestion (self-interference from other devices within the home) are increased. While the increases occur across the board, these increases are particularly problematic in dense living environments such as Multi-Dwelling Units (MDUs), for example apartment complexes.


**Figure 11** shows interference levels in suburban homes vs. urban MDUs as collected on Plume managed Wi-Fi networks. The graphs are histograms of the interference levels. The x-axis is the percentage of airtime consumed by the interference. For example, 33% interference indicates neighbors are consuming 33% of the available airtime, reducing throughput in the home in question by more than a third. The y-axis shows the percentage of homes/apartments in the particular interference bin of the histogram. 100% of homes surveyed are represented at some point in the histogram.





**Figure 11 - Interference Histogram for Suburban Homes and Urban MDUs**

As with previous COVID-19 related issues, the application of sophisticated cloud based processing can mitigate interference. In this case, the key is to consider the MDU as a whole. The algorithm begins with the APs reporting statistics about the neighboring APs they can see, along with interference levels and traffic loads. Based on the lists of neighboring APs, a clustering algorithm, run in the cloud, forms groups of APs that are tightly coupled to each other. The algorithm balances cluster size with the completeness of including all APs that interact with other APs in the cluster. Once the clusters have been formed, the optimizer can select frequency channels and bandwidths so as to maximize the performance of the entire cluster. **Figure 12** shows the result for a sample MDU.




**Figure 12 - Before and After MDU Joint Optimization Interference Histograms**

The upper histogram shows the distribution of homes vs. interference levels before optimizing together as a cluster. The lower histogram shows the distribution after optimizing together as a cluster. The vertical line marks the point at which 33% of airtime in homes is consumed by interference. The joint optimization is able to significantly reduce the number of apartments experiencing interference above 33%, the point at which interference effects become significant.

### 3.4. Cyber Security Attacks and IoT Device Security

It has been hypothesized that criminals would take advantage of the fear, confusion, and increased internet use due to the COVID-19 pandemic. Sadly, it turns out this is true. **Figure 13** shows the increase in various types of threats before and after COVID-19, with several attack types doubling in frequency. In fact, across the period of this study 87% of the homes had some type of cyber security attack.



**Figure 13 - Cyber Security Attacks Before and After COVID-19**


To combat this, a robust cyber security system should be deployed. A system that includes a variety of layered protections will be most effective. Each of the following methods have their own advantages and weaknesses, but when taken together as a whole form an effective defense.

**DNS query checking:** Devices making an outbound connection to the internet often begin that transaction with a Domain Name Service (DNS) query to obtain the IP address for a given domain name. Domain names are generally logical and somewhat understandable, so it is relatively easy to maintain a list of high risk and low risk domains. A disadvantage is that creative cyber criminals avoid the use of domain names by creating viruses that use IP addresses immediately, without a domain name lookup.

**IP address reputation checking:** IP addresses can be checked similarly to domain names. The advantage is that this works well for incoming traffic as well as outgoing traffic, and it can cover cases in which viruses avoid DNS lookups. However, IP addresses change frequently, and maintaining a valid list of acceptable vs. unacceptable IP addresses is difficult.

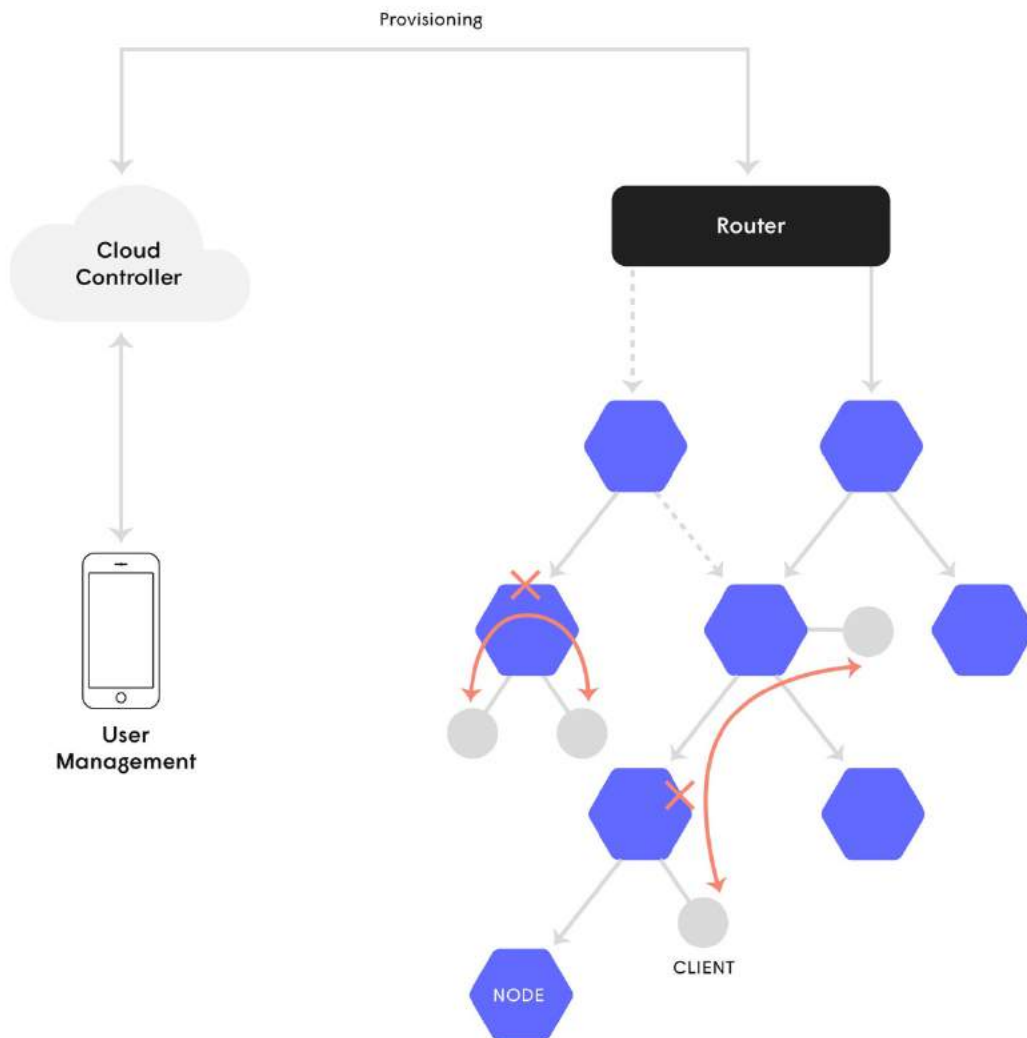
**Port checking and anomaly detection:** Internet connections usually have an associated port number, which is frequently indicative of the application being used. Certain applications, such as ssh, should not be active on certain types of devices. Such basic rules can be augmented by doing machine learning (ML) based anomaly detection. The ML approach can be particularly powerful if data from millions of homes can be aggregated in the cloud, forming a robust dataset for training.

**Tx/Rx Data pattern anomaly detection:** The last line of defense is to observe the basic behavior of the device. For example, **Figure 14** shows a plot of the number of bytes transmitted and received by a Nest Camera in a given minute plotted as a two dimensional scatterplot. Regions that define different modes of operation can be identified, and behavior that does not fall into these regions can be flagged as anomalous, potentially caused by a virus or malware. This technique is particularly helpful for IoT devices, which are generally headless, and can't have aftermarket anti-virus software added. As with all ML based anomaly detection techniques, having a large training set is crucial. Cloud based management systems that aggregate data from millions of homes are beneficial for this approach.



**Figure 14 - Scatterplot of Tx and Rx Bytes per Minute for Nest Camera**

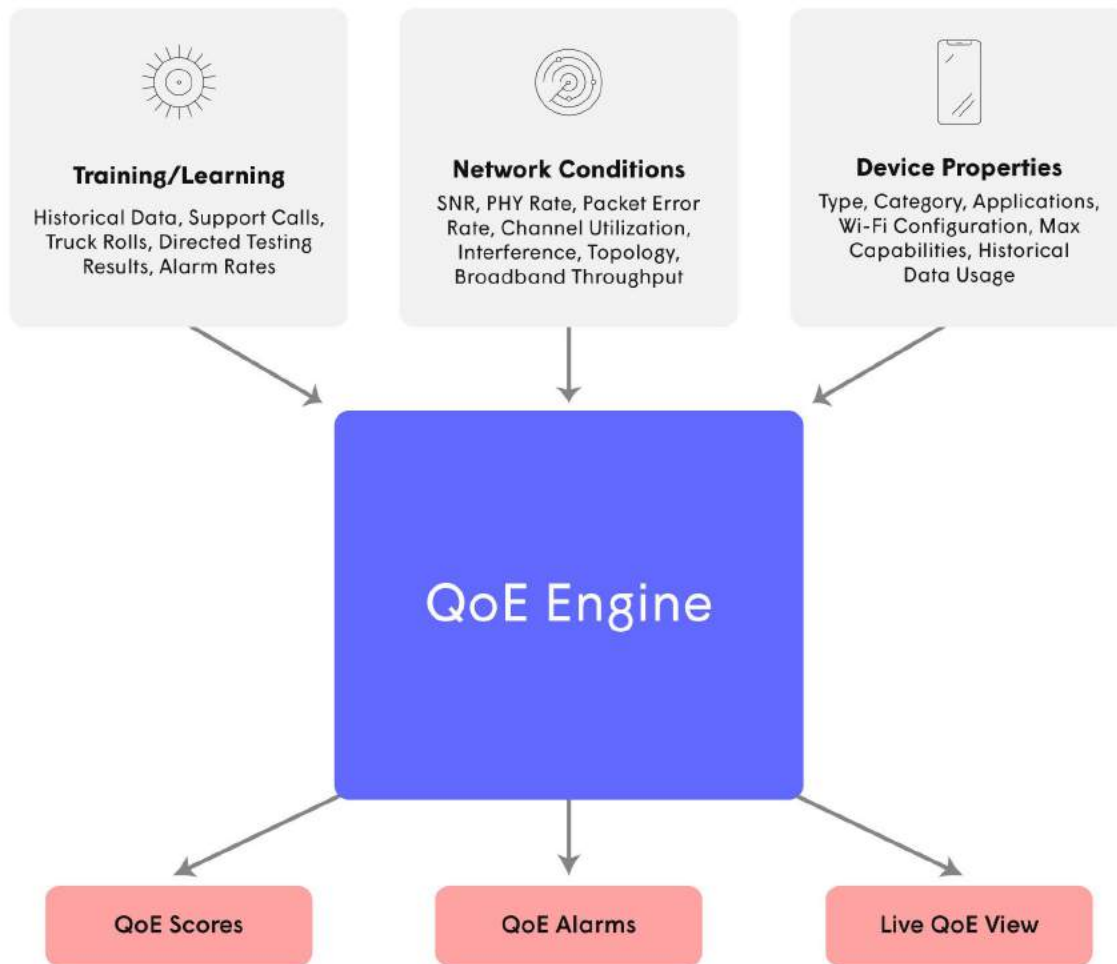
**Cyber Security at every node:** Cyber security measures have traditionally been implemented only at the gateway or main router in a home. However, with the advent of extenders and repeaters, this invites the lateral movement of viruses or malware between devices in the home. Traffic that flows within the home is naturally routed over the shortest path, and this path may not include the gateway or router. And, many mobile devices leave the home and return, sometimes with a new virus on the device. Figure 15 shows the need to detect and stop viruses even for traffic flowing only through a repeater or extender.



**Figure 15 - Lateral Movement of Viruses and Security at Every Node**

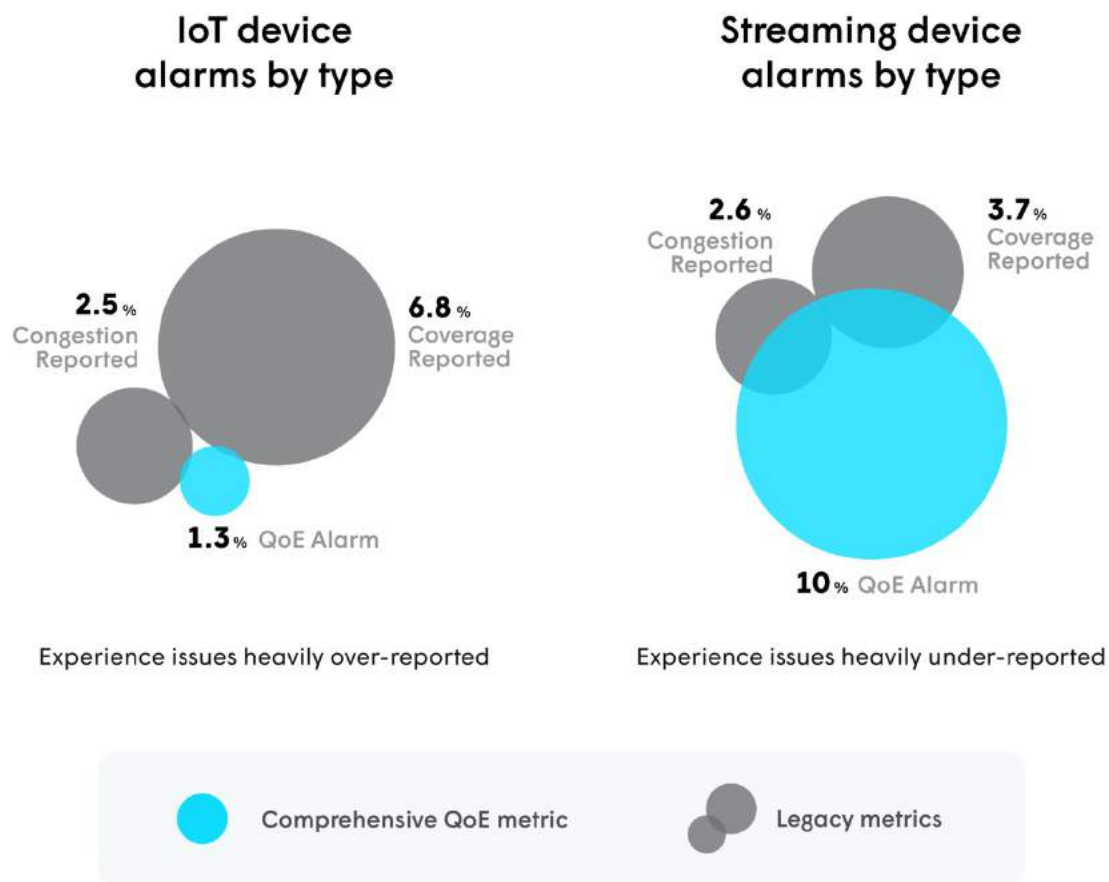
### 3.5. Traditional QoS vs. QoE

As COVID-19 puts more demands on networks, carriers are naturally concerned whether their networks are performing adequately to satisfy their customers. Quality of service (QoS) is the metric that has traditionally been used to indicate this. Typical Wi-Fi QoS metrics factor signal strength, data rate, and congestion/interference, to indicate the quality of the operation of the network. However, consumers are not actually interested in the performance of the network. Consumers are interested in the experience they have using services over those networks. To accurately reflect consumer's satisfaction, Quality of Experience (QoE) is a more effective metric than QoS. QoE starts with similar factors as QoS, but extends that to consider the devices and services that customers are using in a particular home. As an example, a customer with a Wi-Fi thermostat that receives reliable 1Mb/s service will be perfectly happy. On the other hand, a customer with a 4k set top box that receives 20 Mb/s service is likely to be unhappy. Along with factoring the type of device or service in question and their needs, the QoE metric factors far more conditions than just signal strength, data rate, and congestion. **Figure 16** shows the factors considered in a QoE algorithm, and how such a QoE score can be used.



**Figure 16 - QoE Factors and Outputs**

The resulting difference between the QoS and QoE metrics can be compared by using traditional QoS metrics to identify devices that are performing poorly, comparing that set to devices identified using the more accurate QoE metric. **Figure 17** shows the statistics of such a comparison on networks operated by Plume. The Venn diagrams indicate that the different methods identify largely different sets of devices as requiring attention.



\* Based on anonymized data taken from a sample of U.S. households powered by the Plume Cloud

**Figure 17 - Struggling Devices Identified by QoS and QoE**

As could be predicted from the description above, too many IoT devices, which generally have lower networking requirements, tend to be identified as needing attention by traditional QoS approaches. This is significant, as service providers might spend too much money trying to fix those problems with additional repeaters or other steps. Conversely, streaming devices are under identified as needing attention by traditional QoS mechanisms. This also can be costly, as customers experience poor video quality and churn to another service provider.

#### 4. The Effect of COVID-19 on Financials and How to Compensate

Luckily, COVID-19 is not having the devastating financial effect on service providers that it is having on other industries. However, the impact on carriers is still significant. New technologies can help service provider financials in at least two areas: the costs of customer support, and additional revenue from new services.




## 4.1. Proactive Support

As COVID-19 drives heavier use of home networks, and has consumers spending significant time on sensitive applications like teleconferencing, customer support becomes more heavily loaded. An innovative application of machine learning is to predict which customers are most likely to call in the next period of time. Such a capability can be used in a variety of ways:

- Direct proactive maintenance/care
- Send preventive email
- Understand what drives calls

ML cannot predict who is going to call perfectly, but it can do quite well. And, the algorithm can be adjusted to trade precision (the percentage of people identified by the algorithm who really would have called), and recall (the percentage captured on the list of all people who would have called). **Figure 18** shows results that Plume achieved on data from a North American deployment.

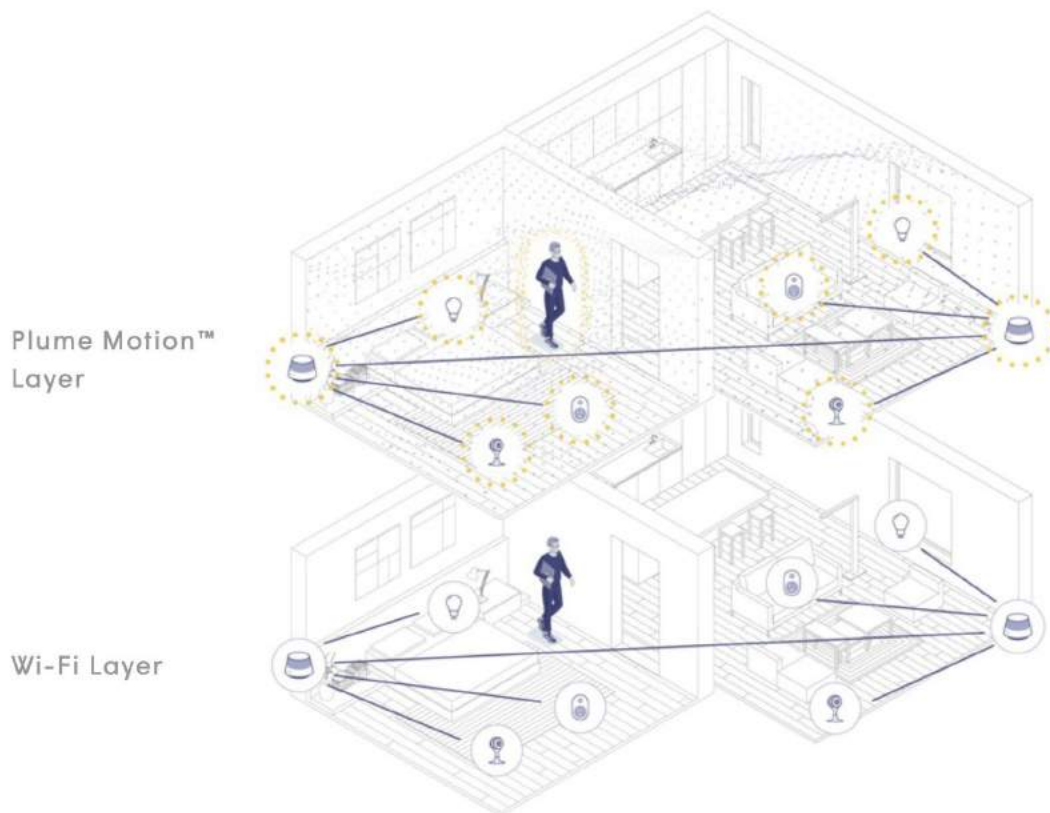


**Figure 18 - Support Call Prediction Precision vs. Recall**

Taking one particular point off the curve, it shows that we can form a list of people expected to call which includes 60% of people who actually will call (the x-axis), and the people on that list will be statistically 10x more likely to call than the population as a whole. The curve shows how these two values can be traded against each other.

## 4.2. Lost Revenue and Compensation with Additional Services

Another approach to make up for profit lost due to COVID-19 is to introduce additional services that users will pay a premium for. An innovative new service that can be offered using the in-home Wi-Fi network is motion detection. Motion in a home can be detected by watching the Wi-Fi signals between devices in a home. These signals will change in time if people are moving in the home due to changes in attenuation and reflection patterns. **Figure 19** illustrates the concept:



**Figure 19 - Wi-Fi Motion Detection Concept**

Using Wi-Fi to detect motion has a number of advantages:

- Far lower privacy concerns than using cameras in the home
- Does not require installation of extra sensors
- Does not require wearing of tags or sensors
- Comes with no hardware cost - can be implemented as a pure software solution on existing Wi-Fi gear

Wi-Fi motion can be used to provide a number of services of interest to the user. Home physical security is an immediate application. When out of the house, the system can be put into a mode where an alarm is sent if motion is sensed in the house. Elder care is another application, based on the reverse concept. In elder care, an alarm is sent if there is *not* motion across the appropriate periods of the day. Home energy management, for example turning on the heat when someone arrives home, is another application. Future versions of the technology may even be able to detect falls, and react quickly enough to control lights as you enter a room.

## 5. Conclusion

COVID-19 has changed the environment for service providers. It has dramatically increased the usage, loads, and sensitive traffic in homes, increasing the difficulty of providing sufficient quality to consumers. And, these changes are likely to persist even after COVID-19 has receded. However, technologies exist that have proven to mitigate many of the issues detailed in this paper. Multiple-AP deployments, with intelligent optimization and throughput optimized steering can ensure sufficient throughput to all locations in the home. Interference in the most challenging environments, multi-dwelling units, can be minimized through cloud based joint optimization. Increased cyber security attacks can be thwarted with machine learning based anomaly detection. And, accurate evaluation of networks requiring attention can be achieved using quality of experience metrics. COVID-19 is also affecting service providers revenue and profit. Artificial intelligence can be applied to perform proactive support, reducing the costs of service calls, truck rolls, and most critically, customer churn. And new technologies such as Wi-Fi motion detection can be used to create new services and new revenue streams.

## Abbreviations

|          |                                     |
|----------|-------------------------------------|
| AP       | access point                        |
| COVID-19 | Coronavirus Disease 2019            |
| DHCP     | dynamic host configuration protocol |
| DNS      | domain name system                  |
| Mb/s     | Megabits per second                 |
| MDU      | multi-dwelling unit                 |
| QoE      | quality of experience               |
| QoS      | quality of service                  |
| UPNP     | universal plug and play             |
| WFH      | working from home                   |

# Power Management on the Generic Access Platform

A Technical Paper prepared for SCTE•ISBE by

**Frank Sandoval**  
Principal Engineer  
Pajarito Technologies, LLC  
Denver CO 80206  
720 338 1988  
Francisrsandoval@gmail.com

# Table of Contents

| <b>Title</b>                                                             | <b>Page Number</b> |
|--------------------------------------------------------------------------|--------------------|
| 1. Introduction .....                                                    | 3                  |
| 2. Power measurement and management on the Generic Access Platform ..... | 3                  |
| 2.1. System Overview.....                                                | 3                  |
| 2.2. Node Manager<->GAP Communications .....                             | 5                  |
| 2.3. APSIS .....                                                         | 6                  |
| 2.4. Power Requirements .....                                            | 7                  |
| 2.5. Power Use Cases.....                                                | 8                  |
| 2.6. Adaptive Power .....                                                | 9                  |
| 2.7. The Application Layer .....                                         | 10                 |
| 2.8. Next Steps.....                                                     | 11                 |
| 3. Conclusion .....                                                      | 12                 |
| Abbreviations.....                                                       | 13                 |
| Bibliography & References .....                                          | 13                 |

## List of Figures

| <b>Title</b>                                       | <b>Page Number</b> |
|----------------------------------------------------|--------------------|
| Figure 1 - GAP .....                               | 4                  |
| Figure 2 - System Overview.....                    | 4                  |
| Figure 3 - YANG snippet.....                       | 6                  |
| Figure 4 – SCTE Energy Pyramid .....               | 8                  |
| Figure 5 - Diurnal Adaptation (illustrative) ..... | 9                  |
| Figure 6 - Full Stack .....                        | 11                 |

# 1. Introduction

Specification of the Generic Access Platform is currently underway (Summer 2020) within SCTE, with broad participation of cable operator and vendor communities. In addition to the requirements for modularity and flexibility driving the effort, the platform will offer unprecedented communications to monitor and control the device and power management is an important feature in those communications.

The cable access network accounts for the vast majority of energy consumed by cable broadband providers, due to the very large number of devices in the outside plant. Measuring energy consumption deep into the network provides a number of business and customer-facing benefits, and the possibility of optimizing electrical usage under certain scenarios offers a glimpse of a more energy efficient, cost efficient, and reliable network.


This paper will briefly introduce GAP and the network configuration (NETCONF) protocol used for its communications, and dive into the SCTE-216 APSIS standard that will provide the power management facilities. APSIS - the Adaptive Power Systems Interface Specification has been developed by the cable operator community and published by the SCTE to provide a comprehensive and flexible data model to represent energy metrics and controls. By adopting a standard data model, costs can be driven out of the data supply chain, and this data can be merged across multiple platforms. A number of APSIS based energy use cases have been developed by the industry, many of which can be directly supported by GAP.

Finally, some suggestions for next steps for the industry will be outlined.

## 2. Power measurement and management on the Generic Access Platform

### 2.1. System Overview

The Generic Access Platform (GAP) is a set of specifications that will be published by SCTE that define interfaces to enable outside plant/access network components from multiple vendors to interoperate to provide a configurable set of functionalities. Unlike a traditional node that embodies a fixed set of features, a GAP chassis can be populated by sub-modules to provide a range of capabilities, and can be updated and reconfigured while in the field. GAP can support any mix of services such as DOCSIS, WiFi, PON, and 5G.

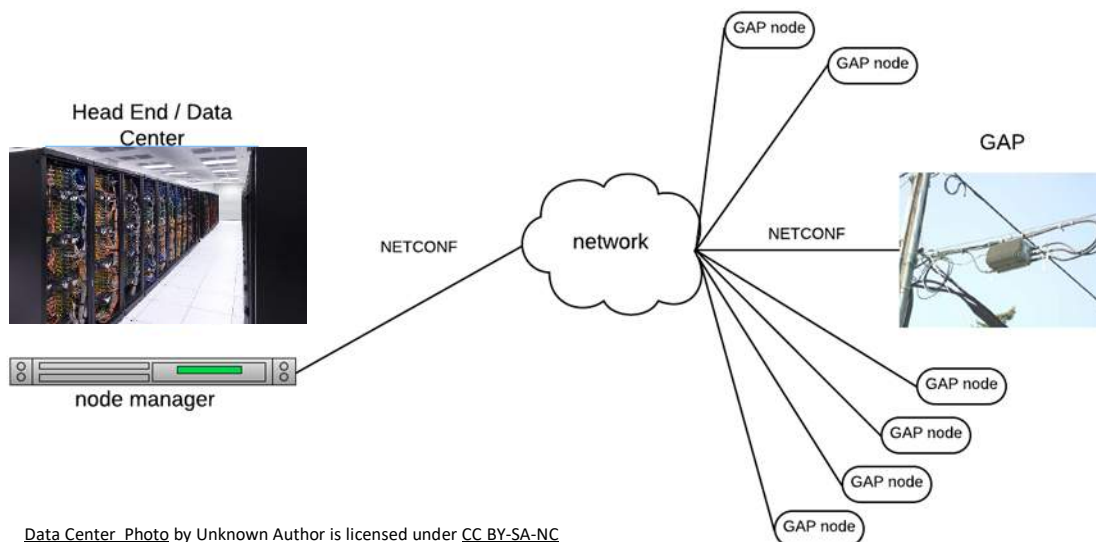


**Figure 1 - GAP**

For more detailed information about GAP, please see the SCTE online description of GAP.

Interoperability demands standardization so that the same type of component supplied by two manufacturers can appear to behave identically to other components in the system. GAP defines how sub-components within a device, such as amplifier, power supply, virtual Cable Modem Termination System (vCMTS) and other components interface with the device chassis. Additionally, GAP defines how the system communicates with an external control component, described as a node manager.

The node manager communicates with the GAP from the head end or other location in the broadband network via the NETCONF protocol.



**Figure 2 - System Overview**

GAP specifications do not define elements outside the scope of the GAP device and its communications with a node manager, but implicit in the design are applications that make use of the node manager to configure and monitor a population of GAP devices. Applications can interoperate with multiple domains because GAP interfaces are well-specified. - For example, an application that listens to the node manager to detect a power outage on a set of GAP device could be leveraged across multiple cable operators. This ‘write-once, run anywhere’ facility can greatly speed innovation and reduce prices where it matters - in the development of the sophisticated logic that exists in applications. Even where applications are developed for a specific provider, the use of standards provides a ‘loose-coupling’ between application logic and the interfaces and data that feed them, reducing development time and maintenance costs. We discuss applications in more detail below.

## **2.2. Node Manager<->GAP Communications**

The GAP communications drafting group has begun listing node level management objects that a node manager can monitor. These are in addition to sub-modules specific telemetry. The initial list includes categories for power, inventory, environmental, security, and other sensors.

There are numerous protocols that could be used for node manager to GAP communications. The communications protocol has to be bi-directional, support a rich data set, be fairly low-latency, and support high volume on the node manager side, as one manager may connect to a high quantity of GAP devices. While traditional legacy SNMP might have been selected, modern protocols are more efficient, secure, and easy to use. Of the possible candidates NETCONF was chosen because it meets all of the requirements and has a healthy ecosystem of tooling available.

The NETCONF protocol has been designed specifically for networking applications and is widely supported in routers, switches, and other networking gear. NETCONF messages are described using a formal DSL (Domain Specific Language) called YANG (Yet Another Next Generation). YANG is a human readable text format that can express hierarchical data trees (similar in some ways to an SNMP MIB), remote procedure calls (RPCs) and notifications.



```

container eoDevices {
 config false;
 list eoDevice {
 config false;
 key "id";
 leaf id {type int32;}
 leaf eocategory {
 config false;
 type enumeration {
 enum PRODUCER {description "energy object category";}
 enum CONSUMER {description "energy object category";}
 enum METER {description "energy object category";}
 enum DISTRIBUTOR {description "energy object category";}
 enum STORE {description "energy object category";}
 }
 description "energy object category";
 }
 }
 uses energyGroup;
 container powerInterfaces {
 list powerInterface {
 key "id";
 leaf id {type int32;}
 uses energyGroup;
 description "collection of power interfaces";
 }
 }
 container components {
 list component {
 key "id";
 leaf id {type int32;}
 uses energyGroup;
 description "collection of components";
 }
 }
}
description "root of enam api";
}

```

**Figure 3 - YANG snippet**

Tooling is available to auto-generate NETCONF interfaces based on YANG, and to generate complementary RESTCONF interfaces. Where NETCONF is a low-level device oriented protocol, with support for connections and transactions, RESTCONF presents a more application friendly HTTP RESTful API.

In typical NETCONF usage, a software defined networking (SDN) ‘controller’ sits between a service provider’s business applications and networking elements; presenting an easy to use RESTCONF API “northbound” to applications and handling the complexities of interacting with devices via NETCONF on “southbound” interfaces. In the GAP scenario, a node manager embodies the functions of an SDN controller.

### **2.3. APSIS**

One of the initial efforts within the SCTE Energy 2020 initiative was to define software interfaces to measure and manage power. A working group was formed to develop the Adaptive Power Systems Interface Specification, or APSIS. Since engineers often like to borrow the work of others rather than build from scratch, a survey of existing power related standards was conducted. The Internet Engineering Task Force (IETF) Energy Management (EMAN) framework was selected as a basis from which APSIS could evolve. Originally defined as a collection of SNMP (Simple Network Management Protocol) MIBs (Management Information Base) contributed by Cisco’s EnergyWise team, EMAN provides a

comprehensive and flexible data model for characterizing the power consumption, and production, of any sort of system. The EMAN structure applies well to the cable domain as it can describe a device with any number and configuration of sub-components (like GAP) and can be easily extended to accommodate special cases if the need were to arise. While the APSIS data model is fairly large, specific use cases need only utilize the portions that are relevant in their context.

After selecting IETF EMAN, the APSIS team developed a high-level, protocol independent Information Model and contributed it back to the IETF. The Information Model describes, in Universal Modeling Language (UML), the same data as encoded in the original EMAN MIBS, but in a way that facilitates additional protocol ‘bindings’. With this mechanism, teams can develop NETCONF, IPDR, gNMI, or any other style of interface to suit their domain yet be confident that the resulting data can be merged with data sourced using another APSIS compliant protocol.

A protocol independent data model is extremely valuable as we consider the job of processing data at higher logical layers. Consider a predictive modeling application that consumes power data from a wide variety of devices to correlate variables in order to anticipate a service outage - perhaps an unusual oscillation of power quality is a predictor of service failure. If power quality measures are being collected from the widest possible population of sensors, perhaps some using SNMP, some NETCONF off of GAP, some using something else; the modeling application can only reliably utilize that data which is semantically identical, even where the on-the-wire syntax may differ. Where data sources use different models to describe similar things the data quality of the merged set can be unreliable. Data scientists can tell us that even where two differing data source formats use a common key name there is no guarantee the associated values mean the same thing.

By referencing APSIS in the GAP specifications, a reliable data model is established on which to build data processing applications. APSIS can further be adopted by other platforms, either as a native format directly supported by a device, or as a target format for which a defined mapping exists from a non-APSYS or legacy data model in the form of the APSIS model. Another opportunity for expanding the footprint of native APSIS support could be through upcoming revisions to the SCTE power supply/transponder interface specification, SCTE 25-3.

## **2.4. Power Requirements**

Among the node manager to GAP communications requirements identified by the GAP working group, a number of power related requirements have been defined, including input voltage, power efficiency, line usage, and power per line, to name a few.

Power consumption is important to proper operations of the access network and it represents a vast majority of electricity costs for a cable provider. As part of the SCTE Energy 2020 program, the cable energy pyramid was published to illustrate the relative energy utilization across portions of a cable operator’s footprint.



**Figure 4 – SCTE Energy Pyramid**

For a large MSO, the access network alone can consume hundreds of millions of electricity dollars in a year.

## 2.5. Power Use Cases

To this point, little has been done to systematically collect, analyze, or act upon energy data emanating from the access network, outside the scope of individual efforts within some MSOs. By including standardized energy metering and controls into the GAP, MSOs and partners may more cost-effectively build the tooling necessary to fully leverage such data.

A substantial number of use cases pertaining to powering the access network have been articulated by the industry, published as SCTE-245 'Use Cases for Adaptive Power Using APSIS'. Use cases have been grouped under several categories, including:

- Measurement
- Adaptation
- Demand Response
- Energy Supply monitoring
- Energy Services


As Lord Kelvin stated, "If you cannot measure it, you cannot improve it." Measurement of energy usage is a fundamental use case. Gaining near-real-time visibility into the sources powering the access network and the power state of components therein supports any number of valuable applications, including at least:

- **Detecting grid power availability.** The cable plant overlays the utility grid and can report with greater resolution the state of power outages. The CableLabs Gridmetrics project is working across the industry to support this and related use cases.
- **Assessing grid power quality.** Grid power is not simply off or on, fluctuations in voltage can impact system performance, reliability, and availability, and can have a dramatic impact on the useful life of equipment.

Please refer to SCTE-245 for details on these and many other important use cases.

## 2.6. Adaptive Power

The seed of the APSIS effort was the observation that while cable service demand swings widely between prime-time peak hours and the middle of the night, there is very little corresponding fluctuation in energy usage by the network. The service demand curve is very similar to what's been called the 'duck curve' in the utility industry since the shape of electrical load over the course of a day resembles the silhouette of a duck's back. As cable technologists, the fact that energy consumption remains high when service delivery plunges just feels like a system that is not optimally designed - it's a bit like leaving the lights on when you leave a room at night. The use case of adapting network behavior to correlate energy consumption to the stable and predictable daily service demand oscillation has been labelled 'diurnal adaptation'.



**Figure 5 - Diurnal Adaptation (illustrative)**

Figure 5 attempts to illustrate that the combined blue and grey areas are the flat, unoptimized energy consumption, and the blue is optimized by the effects of diurnal adaptation. The curves do not represent measured data but are simply illustrative of the phenomenon of daily service demand fluctuation and associated energy curves.

Research by vendors, including ARRIS (now CommScope), WES.NET, and Concurrent, demonstrated at earlier SCTE Cable-Tec Expo events, indicate that by carefully managing network resources, daily energy consumption by some networking elements, such as a Converged Cable Access Platform (CCAP), can be reduced by over %15. In the case of CCAP, as data throughput drops below defined thresholds, flows can be remapped to consolidate traffic onto fewer output ports, freeing up line cards that can then be temporarily placed into low power states.

A proof of concept conducted by Comcast with support from ARRIS, also discussed in a previous Cable-Tec Expo, demonstrated the use of APSIS in managing a CCAP to simulate diurnal adaption and confirm the results of earlier prototypes. The open-source OpenDaylight (ODL) SDN controller was selected to act as the control plane, playing a role similar to that of the node manager in the GAP scenario. An EMAN 'plug-in' was developed for ODL to present APSIS as a standardized Northbound API to an

energy management application, and a Southbound ‘adapter’ was developed to integrate with the command line interface (CLI) of the CCAP. A time-lapsed simulation of daily data traffic was fed into a lab installation of a CCAP using the iperf tool, while the energy management application monitored the CCAP to detect the data throughput. When throughput crossed a defined threshold, the application sent commands, via APSIS APIs, through the controller to the device to remap flows and power up/down line cards, as the volume of traffic dictated.

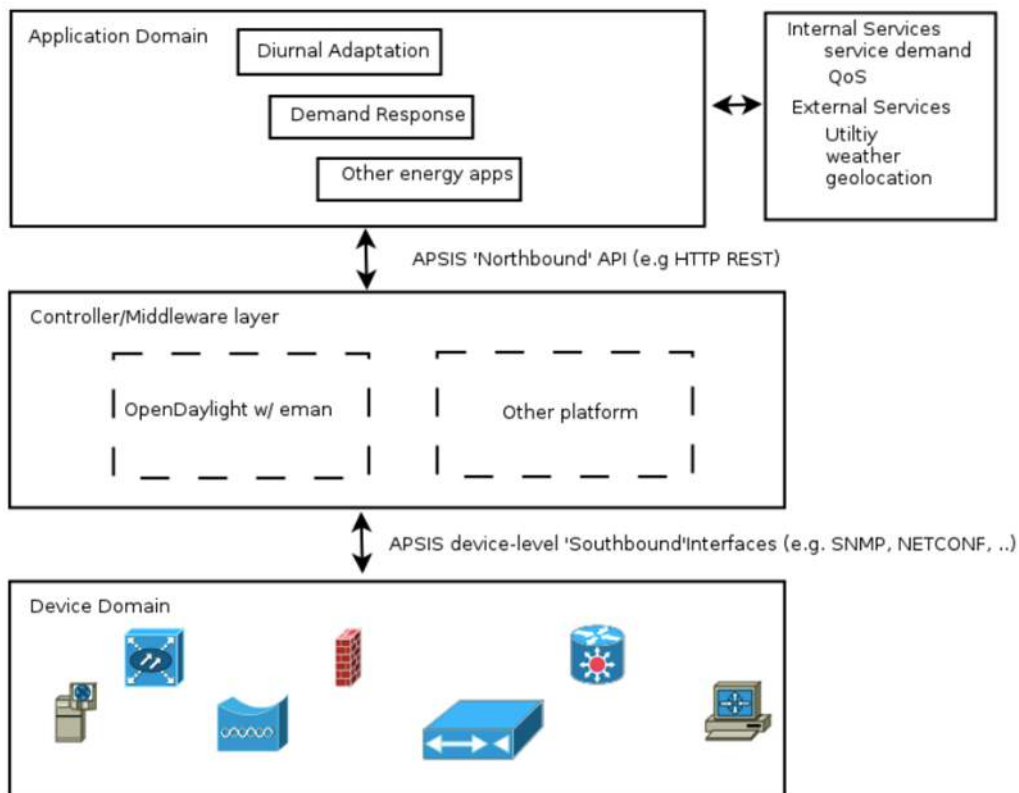
Adaptive control of the GAP could take several forms, including attenuation of the bias current driving the downstream radio frequency (RF) signal. The strength of the baseline bias signal, on to which service information is encoded and carried, can be correlated to the amount of information that needs to be carried. As data volumes increase and decrease to serve shifting demand, the strength of the baseline bias current could correspondingly attenuate. The strength of the signal is a function of the power used to generate it; therefore, a diminished bias current consumes less electricity. In a generic access platform, bias current might be controlled by the signal generator in response to the current service demand, or an external application that has access to other data sources and predicative models could augment this logic to provide a more timely and accurate control algorithm.

In a production environment, such real-time manipulation of the network could only be attempted with sophisticated control mechanisms. A production diurnal controller application would not only monitor data throughput but would factor in a variety of other metrics to maintain consistent quality of service and preserve the custom experience. An application could monitor customer experience metrics such as packet loss and jitter, and predicative modeling of anticipated data rates generated from Machine Learning agents. In addition, weather events, social events like sports, or other phenomenon that might perturb the historical norm could be factored in. The application would have to implement automatic and manual controls to back-off its optimization operations in order to ensure continuity of service. While this seems futuristic today, as network operations have been successful in maintaining their excellent service availability record by minimizing risks and limiting variables, we can anticipate that as software controls become more sophisticated they will be able to lower costs and increase customer satisfaction while addressing the risks of unintended consequences.

Adaptive control is a type of ‘closed-loop automation’; a topic of intense focus within the telecommunications world at large. The dream of self-configuring and self-healing networks is in development today.

## **2.7. The Application Layer**

Interface standards provide the plumbing through which data and commands can flow, but interface standards are valuable only in what business logic they support. When we include applications into our system view, we get a three-tiered ‘full stack’ model: devices at the bottom, node manager/SDN controller in the middle, and applications at the top. GAP is a special case of a more general data ingest/processing/business logic framework, as power related data is just one of a number of data streams that can and should be handled in the same general way. Another way of describing the general case is standards and processing system are the domain of data engineering, and the transformation of data into business intelligence is the domain of data science.



**Figure 6 - Full Stack**

Figure 6 illustrates a general model for a device/controller/application stack. GAP is a specific implementation of this general model, in which the ‘Controller’ in the node manager and the device domain is limited to GAP.

A healthy ecosystem utilizes standardized data flowing into cable operator data processing systems feeding an ever-growing number of applications. To foster innovation at the application layer, apps must utilize known interfaces and data models - otherwise developers will spend valuable time integrating with bespoke data sources rather than focused on creating valuable business logic.

Cable operators will use their in-house teams to develop some applications, but may also partner with application vendors (exporting data under strict privacy and security controls to authorized and authenticated partners), among whom might include value-added analytics services, utilities, industry consortia, researchers, and any other partner an operator wishes to engage.

## 2.8. Next Steps

Writing and publishing a standard is only worthwhile if it is adopted by multiple industry actors and lead to the creation of business value as previously mentioned in section 2.7. We can greatly speed adoption by performing activities above and beyond simply writing things down on a piece of (virtual) paper and posting it on the web.

The Java Community Process (JCP) provides a great example. A new Java API standard is not considered complete until three deliverables are made available: a specification, a Reference Implementation (RI), and a Technology Compatibility Kit (TCK), or test kit.

We might collaborate across the industry to develop a simple prototype of a node manager linked to a GAP simulator to serve as a first-generation Reference Implementation and make API calls to the RI to serve as the beginnings of a test suite.

Because GAP is using NETCONF, our initial prototyping would be quite easy. An NETCONF open-source SDN controller, such as OpenDaylight, could act as a Node Manager, and one of several NETCONF server simulators could present itself as a GAP, by serving up GAP YANG models. A couple of net-conf simulators are: ODL netconf testtool and ntsim.

With a working environment programmers and integrators can then begin to develop application prototypes to learn how best to process, analyze, and act upon GAP data, regardless of whether power related data or other GAP data. The value of an accurate simulation environment cannot be understated for application development - the key is having a system that supports very tight incremental code changes, e.g. tweak/modify a line of code, test it, repeat forever.

The GAP RI could be run on a laptop or could scale up in the cloud. If someone were to donate cloud resources, like a small AWS environment, the industry could share a ‘GAP lab’ to co-develop the RI, tests, and applications.

Finally, the adoption of APSIS by every platform that contributes to power monitoring or controls will drive costs out of developing data pipelines and applications that make use of power data. Think of plumbing a house. If: the sizes of pipes and fittings were all different, made up on the fly, changing from manufacture to manufacturer, and even changing between years and product lines from a single supplier, all of these factors could cause enormous costs and complexity to the modern day homeowner. This indeed was the case in the early days on modern plumbing. Let’s not waste time and money pursuing non-standard solutions to data formats, including power data.

### **3. Conclusion**

The Generic Access Platform promises an important evolution in access network technology by establishing a modular and configurable node architecture. Among the many benefits to operators is much improved visibility and control, including power measurement and management. GAP should incorporate the NETCONF protocol for communications with an upstream node manager, driven by well specified data models in the YANG format. For the power components of the communications platform, the SCTE APSIS specification provides an excellent solution as it defines a multi-protocol information model and a YANG binding to provide a comprehensive power measurement and management interface.

As a specification of GAP continues, we call for multi-party collaboration to develop a lightweight prototype implementation to ‘burn in’ the spec and identify gaps or mistakes in the written specification, and to provide an application development platform to generate tests and to research commercially valuable applications that process GAP data and support business value.

Finally, we encourage the adoption of APSIS for all power related interfaces and applications throughout the industry. There is little justification to re-invent a wheel and generating data in differing and potentially incompatible formats will only limit opportunities in deriving the utmost value in power data. Data, like social networks, can exploit ‘network effects’ to increase in value as they are merged and correlated with other data.

## Abbreviations

|          |                                                 |
|----------|-------------------------------------------------|
| 5G       | Fifth-Generation cellular wireless              |
| API      | Application Programming Interface               |
| APSYS    | Adaptive Power Systems Interface Specification  |
| AWS      | Amazon Web Services                             |
| CCAP     | Converged Cable Access Platform                 |
| CLI      | Command Line Interface                          |
| DOCSIS   | Data Over Cable Service Interface Specification |
| DSL      | Domain Specific Language                        |
| EMAN     | Energy Management                               |
| HTTP     | HyperText Transfer Protocol                     |
| GAP      | Generic access platform                         |
| IETF     | Internet Engineering Task Force                 |
| ISBE     | International Society of Broadband Experts      |
| JCP      | Java Community Process                          |
| MIB      | Management Information Base                     |
| MSO      | Multiple System Operator                        |
| NETCONF  | Network Configuration                           |
| ODL      | OpenDaylight                                    |
| PON      | Passive Optical Network                         |
| RI       | Reference Implementation                        |
| RESTCONF | RESTful Configuration                           |
| RF       | radio frequency                                 |
| RI       | Reference Implementation                        |
| RPC      | Remote Procedure Call                           |
| SCTE     | Society of Cable Telecommunications Engineers   |
| SDN      | Software Defined Networking                     |
| SNMP     | Simple Network Management Protocol              |
| TCK      | Technology Compatibility Kit                    |
| UML      | Universal Modeling Language                     |
| vCMTS    | virtual Cable Modem Termination System          |
| WiFi     | Wireless Fidelity                               |
| YANG     | Yet Another Next Generation                     |

## Bibliography & References

SCTE Standards Library: <https://www.scte.org/download-scte-isbe-standards/>

SCTE 216: Adaptive Power Systems Interface Specification

SCTE 237: Implementation Steps for Adaptive Power Systems Interface Specification (APSYS)

SCTE 245: Use Cases for Adaptive Power Using APSIS

SCTE Energy 2020; <https://www.scte.org/energy-2020-powering-cables-success/>

RFC7326: Energy Management Framework; Internet Engineering Task Force



RFC6241: NETCONF Network Configuration protocol; Internet Engineering Task Force

RFC7950: YANG Network Configuration protocol; Internet Engineering Task Force

SCTE description of GAP; <https://www.scte.org/generic-access-platform/>

ODL netconf testtool; <https://docs.opendaylight.org/projects/netconf/en/latest/testtool.html>

ntsim; <https://github.com/Melacon/ntsim/tree/master/ntsimulator>

CableLabs Gridmetrics; <https://gridmetrics.io/>

Arris APSIS demo; <https://www.cablefax.com/uncategorized/arris-energy-saving-idea>

Cisco EnergyWise;

[https://www.webopedia.com/TERM/C/Cisco\\_EnergyWise.html#:~:text=Cisco%20EnergyWise%20is%20a%20Green,between%20network%20devices%20and%20endpoints.](https://www.webopedia.com/TERM/C/Cisco_EnergyWise.html#:~:text=Cisco%20EnergyWise%20is%20a%20Green,between%20network%20devices%20and%20endpoints.)

# **Framework for Convergence of Services on The MSO Network**

## **Using the Principles of Network Slicing**

A Technical Paper prepared for SCTE•ISBE by

**Fernando X. Villarruel**

Chief Architect

Ciena

1185 Sanctuary Pkwy, Alpharetta GA, 30009

[fvillarr@ciena.com](mailto:fvillarr@ciena.com)

**David Reale**

Network Architect

Ciena

1185 Sanctuary Pkwy, Alpharetta GA, 30009 [dreale@ciena.com](mailto:dreale@ciena.com)

# Table of Contents

| Title                                      | Page Number |
|--------------------------------------------|-------------|
| 1. Introduction.....                       | 3           |
| 2. Opportunity Statement.....              | 3           |
| 2.1. Network as a Service.....             | 4           |
| 3. Solution Statement.....                 | 4           |
| 3.1. Convergence View .....                | 4           |
| 4. Convergence Drivers .....               | 6           |
| 4.1. Distributed Access Architectures..... | 6           |
| 4.2. Cloud Native Service Cores .....      | 6           |
| 5. 5G and Network Slicing.....             | 8           |
| 5.1. 5G Functions and Descriptions .....   | 9           |
| 5.2. What is Network Slicing .....         | 9           |
| 5.2.1. Service Requirements.....           | 10          |
| 5.2.2. Operation Requirements .....        | 10          |
| 5.3. Service Convergence and Slicing ..... | 10          |
| 6. MSO and 5G coexistence .....            | 12          |
| 7. Slicing methods.....                    | 13          |
| 7.1. Soft Slicing .....                    | 13          |
| 7.2. Hard Slicing.....                     | 14          |
| 7.3. Hard And Soft Slicing.....            | 15          |
| 8. Network Slice Lifecycle .....           | 16          |
| 9. Industry Recommendations .....          | 16          |
| 10. Conclusion .....                       | 17          |
| Abbreviations.....                         | 17          |
| Bibliography .....                         | 18          |

## List of Figures

| Title                                                                   | Page Number |
|-------------------------------------------------------------------------|-------------|
| Figure 1 - Convergence of services for MSOs.....                        | 5           |
| Figure 2 - Evolution to Distributed Access Architecture, DAA .....      | 6           |
| Figure 3 - Cloud Native Framework.....                                  | 7           |
| Figure 4 - Typical Cloud Native Core Deployment.....                    | 7           |
| Figure 5 - 5G Usage Scenarios, (SANTO, 2017).....                       | 8           |
| Figure 6 - Cloud Native 5G Core.....                                    | 8           |
| Figure 7 - End to end network slicing representation .....              | 10          |
| Figure 8 - Domain and Slices Relationship.....                          | 11          |
| Figure 9 - Service Footprint .....                                      | 11          |
| Figure 10 - CMTS / 5G cores coexistence.....                            | 12          |
| Figure 11 - General QoS mapping CMTS and 5G core (CableLabs, 2020)..... | 13          |
| Figure 12 - Soft Slicing Example for Convergence .....                  | 14          |
| Figure 13 - Hard Slicing Example for Convergence .....                  | 14          |
| Figure 14 - Flexible Ethernet Operational Definition .....              | 15          |
| Figure 15 - Hard and Soft Slicing Example.....                          | 15          |

## 1. Introduction

The digitization of the cable access network enables the promise of end-to-end network convergence for different lines of service. The promise, however, does not come with a framework and this leads to trepidation on how to proceed. A framework for service convergence must recognize three key principles. The first is that lines of services can have unique prioritization, throughput and latency requirements that need to be met. The second is that within a service there will be distinctions of endpoint types and applications that will eventually need their own unique treatment or policy through the network. The third is that the principles mentioned above are transitional over the lifecycle of the service and automation mechanisms to adapt and coexist are necessary to maintain viability over the long term.

In this paper we propose a framework for service convergence using network slicing for MSO networks. We review the network slicing principles for 5G and point to possible analogies that aid in developing a framework for MSO slices including residential services, business services, and mobile services. We cover the concept of network slicing functions which organize and partition network resources available to each service. We describe hard and soft slicing mechanisms, the implementation for slice-aware logical networks and the functionality necessary to maintain end-to-end slice visibility and usability over the lifecycle. We also provide several industry recommendations useful for a converged service environment utilizing network slicing, such as open interfaces for core functions, QoS implementation in packet networks, listing of slice expectations, usability of hard and soft slicing.

## 2. Opportunity Statement

The state of the cable industry has top-of-mind investments and revenues from mobile services, the growth of enterprise business services, and maintaining a robust residential high-speed internet and video service. There are multiple MSOs worldwide that already own and manage all three services, and for the last decade in the United States there have been coalitions formed with wireless providers to collaborate in transmission and management of wireline and wireless services. What has been noted, however, is that this is a financially complicated situation when done piecemeal, and that profitability has better odds coming from infrastructure-based network convergence. (BAUMGARTNER, 2019)

Enterprise business services on the other hand has been a solid growth engine for MSOs in recent years, becoming a solid primary connectivity option for enterprises. As a metric, in the past several years cable providers continue to make their way up in the Vertical Systems Group Carrier Managed SD-WAN Services Leader Board with a placement recently in the top seven (Vertical Systems Group, 2020). But one of the drawbacks to even more growth has been time to deployment and SLA enforcement partly due to a dependency on existing methods of non-automated bring up of circuits and separate networks including operation and engineering. Thus, the enterprise network depends on evolution to automation, for itself, in the context of automation being also useful for mobile and residential services.

The residential network continues to have solid growth, with revenues for high speed internet growing at 5-10% year-over-year (Comcast, 2020), (Charter, 2020). The status of the residential network exploring convergence with mobile services, and no convergence with business services is understandable due to historical precedents, but there are several dynamics at work in the evolution of the residential network that allows us to take a renewed look at converging multiple services on common network infrastructure. One is the nature of Distributed Access Architectures (DAA), where the legacy residential cable plant, both fiber and coaxial plants have used analog technologies for distribution, but are now transitioning the fiber portion to digital Ethernet and IP, thus all networks will now approximate the same format for transmission. Second is the evolution of DOCSIS and Video to cloud native cores. Now DOCSIS, Video,

Mobile and general broadband network gateways that handle subscriber management will have a deployment path to use containerized computation schemes with generic hardware support. And third is the large expansion of endpoints that need connectivity for 5G and DAA and enterprise 1/10GbE fiber services which could benefit directly from the very broad coverage of HFC and its availability of electrical power sources to distribute cells, Remote-PHY Devices (RPD) / Remote-MACPHY Devices (RMD), and aggregation points for CPEs (Chamberlain, 2018). From this perspective, the industry is ripe for full infrastructure convergence for the services mentioned above leveraging investments in an optimized and streamline manner.

From the user perspective convergence assists in typical usage methods. Take for example the smart phone user maintaining a session at a home, with Wi-Fi enabled backhaul, proceeding to a car driving down the road all on seamless session. Or the parent working from home at the same time and on the same high-speed internet that provides over-the-top video entertainment for the kids. These examples include seamless handoffs between different service types with different usage priorities.

## **2.1. Network as a Service**

The drivers for network convergence not only include reduction of OPEX and improved general efficiency, it also includes the capability to add new revenue streams. Network convergence organizes the qualities of the network such that the network itself is an asset to sell. When the network understands the user transmission profile and its tendencies, then the path is set for creating a robust platform for network as a service (NaaS), where the operator offers highly customizable Virtual Network Functions (VNF) for customers who need a broad range of network capabilities (Hodges, 2019). A complete description of NaaS capability is left for another work, but necessary to mention in the context of a network slicing byproduct.

## **3. Solution Statement**

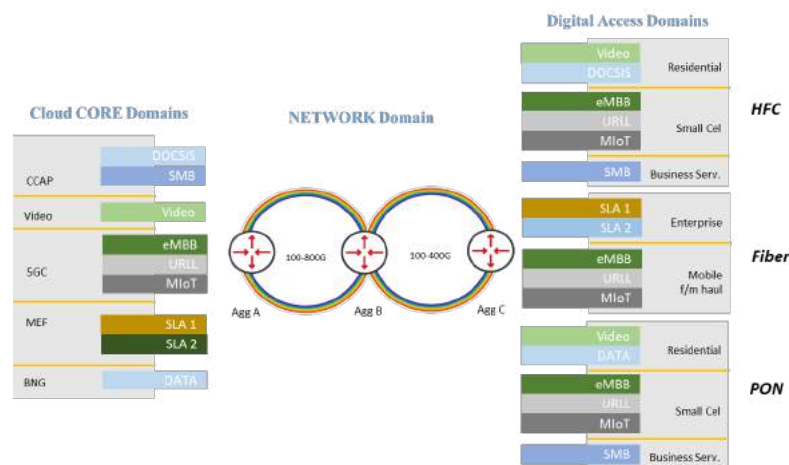
The nemesis to converging residential, private line business services, and mobile services on common network infrastructure boils down to a possible acrimonious sharing of resources. After all, the networks mentioned above have grown up independently, with full control of their infrastructure, customer base, their priority structure and SLA compliance. These are valid concerns of a converged environment that merit a comprehensive solution. This paper proposes a method for developing a non-acrimonious solution for network and service convergence where the sharing of resources is possible. To do this we borrow from the principles of network slicing as recently defined for 5G. Naturally, an implementation of network slicing for the MSO will have its own idiosyncrasies consistent with our needs and expectations.

Before we discuss network slicing, we review a few technical principles that are key to understanding network slicing.

### **3.1. Convergence View**

It is useful to catalog the extent of the convergence possibilities. Figure 1 provides a list of the different type of physical networks, services, and subservices that would be part of a converged end to end network. We list three access network types, the HFC plant, a point to point fiber plant, and a passive optical network. All these access types would converge on the same aggregation point C, at a hub or in the field, such as an unmanned cabinet or a strand-mounted device. Aggregation at point C could be done in fiber, per wavelength, in straight forward Time Division Multiplexing (TDM) using Optical Transport Network (OTN) or Flex-Ethernet (FlexE) framing, in layer 2 switching using Ethernet, or layer 3 routing using Internet Protocols. Northbound the aggregated signal could range from 10 Gbps to 400 Gbps

depending on optical technology and desired level of signal concurrency (Villarruel, 2018). The aggregated signal would then typically terminate in another aggregator, point B, whose job is to add/drop signals locally if a service core is present, or further aggregate signals onto a metro type network with signal rates in the 100 – 800 Gbps range. The technologies northbound are typically IP with embedded high throughput optics in the range of 100-400 Gbps or a more advanced photonic layer in the range of 100-800 Gbps. Aggregation point A is then the termination point for the signals that originated in the access. Here they are evaluated in accordance to session and subscriber management at independent services cores. The Converged Access Platform (CCAP) takes care of the DOCSIS signaling and represents in this case other auxiliary cores of other residential type functions, such as out of band set top box signaling and test and measurement. The video core cares for the distribution of MPEG sessions, broadcast and narrowcast. The 5G core cares for the mobility signaling. With the understanding that there is an ongoing evolution from LTE to 5G core, we shortcut to 5G for sake of expediency in this discussion.) The Metro Ethernet Forum (MEF) core stands for the session management for enterprise and applicable business customers. The Broadband Network Gateway (BNG) is the session and subscriber manager for the data subscribers of the Passive Optical Network (PON) network.



**Figure 1 - Convergence of services for MSOs**

Within each access network there are different types of services. Figure 1 shows a non-exhaustive list of service types. In the HFC access, for example, there are wireless small cells, which are also expected to use DOCSIS as backhaul (Andreoli-Fang, 2019), along with business customers. For Fiber access there are enterprise customers and mobile front / midhaul or backhaul customers. For PON access there can be residential, small cell and business services customers.

Figure 1 also shows that within each service vertical there are subservices, this is also a non-exhaustive list. These are unique user applications that have specific demands of their parent network consistent with the end user expectations. Within HFC/residential there is Video and DOCSIS data. In the HFC/small cell domain we list enhanced Mobile Broadband (eMBB), Ultra-Reliable Low-Latency Communication (URLLC), and Massive IoT (MIoT), which are specifically defined sub-service types for 5G. For Fiber/enterprise there are SLAs and in both PON and HFC there are SMB Business Services. Note that the same type of subservice could live within different access domains.

Now these different types of customers/services are managed by their respective core domain. In practice, and for the foreseeable future, we can assume that even though there is a transition to cloud-based cores, the cores would remain independent, i.e. not sharing compute or memory resources. Figure

1 however presents the cores as converged in the minimum sharing resources, (while keeping functional independence), but at some point, in time possibly sharing session and subscriber management functions.

Figure 1 also shows that subservice types have a unique termination point on the core, independent of access type they originate from. This is also a forward-looking proposition, as we recognize that each subservice now subtends to its own specific service core.

## 4. Convergence Drivers

### 4.1. Distributed Access Architectures

DAA is a technical driver for convergence. This is a topic that has received much due attention in the past few years and here we refer to work that has already been done. DAA is driven by the evolution of the residential network from analog fiber to digital fiber by extracting the physical RF layer from CMTS or effectively extending the Quadrature Amplitude Modulation (QAM) modulation platforms and positioning them in a separate location, where separate can mean a different shelf, in a different hub, or at the end of the deeper fiber additions in the outside plant, in a street cabinet or on strand mounted node like platforms. DAA technologies create an Ethernet / IP network where several aggregation points are needed to distribute or collect signaling. DAA because of its standards based digital transmission, is a natural platform for usage of other access endpoints that coincide using ethernet or IP access signaling (Villarruel, 2014). Ethernet point to point signals for business services apply, and so do recent pluggable OLT technologies that are granular OLTs, sprouting PON networks from any given 10GbE switch port (Villarruel, 2015). Figure 2 shows the evolution of the residential access plant, starting with analog based fiber then adding digital endpoints, followed by aggregation in the field and ultimately convergence of other services. On the core side it begins with legacy video and CMTS platforms evolving to cloud native cores along with a possible centralization of cores from smaller hubs to larger hubs or from hubs to headends, maximizing efficiency of compute. Not shown but also part of the DAA discussion is partitioning of software functions so not all of the DOCSIS stack resides in one location, this for example is the effort being done in the Flexible MAC Architecture (FMA) working group at CableLabs (BTR , 2019).

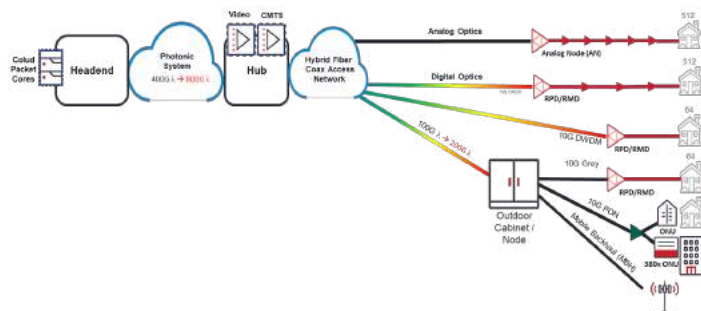



Figure 2 - Evolution to Distributed Access Architecture, DAA

### 4.2. Cloud Native Service Cores

The separation of specific PHY implementations in the CMTS has allowed a rethinking of its software architecture. Effectively virtualizing the CMTS by the decoupling of the software stack from vendor specific hardware. Beyond just virtualizing however there is a trend towards even more software flexibility by implementing a cloud native architecture where software is broken up by functions, and

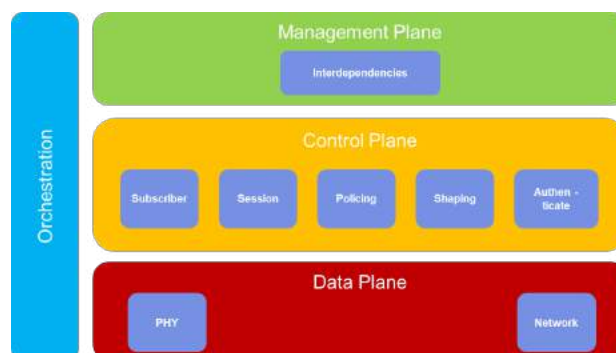
these functions are presented as independent containers that run on a container platform and are orchestrated by a framework that facilitates their interrelation and thus present a complete service solution. Interestingly, these container frameworks are built with open source tools so both the hardware and software infrastructure for service cores is now open, in principle. Figure 3 is a simple view of popular cloud native implementations, where Linux is the operating system, Docker is the container platform and Kubernetes is the orchestrator, all of which were derived in open source communities.



**Figure 3 - Cloud Native Framework**

The nature of containerization also allows for flexibility in geographical placement of containers, allowing for certain core functions to be closer to the user as needed. This is the science of edge compute, derived from cloud native architectures. This topic is beyond the scope of this paper, but it is a discussion that comes up in advanced implementations of distributed service cores and network slicing.

Also, of interest is the view of cloud native cores in deployment as shown in Figure 4. We see the system as a whole is broken up into functionalities. There is the Data Plane, the part of the software that processes and executes on data transmission requests, this notably includes the PHY layers and forwarding layers. There is the Control Plane that creates the configuration environment for the data. This includes things like path controls for a switched or routed network, session and subscriber control for the end user, bandwidth shaping for the network and authentication of endpoints. Beyond the data and control plane there is typically a Management Plane, which can also be part of the control plane, but here we call it out separately to make a point. This management plane facilitates the interdependencies of the control plane within itself and to the physical network items. In scale this typically cannot be done with simple network management system tools, and so an overall orchestration system is implemented. This orchestration system can have read write capabilities for all hardware and software elements and can provide a stateful view of the overall network to northbound operational and billing systems. This generalized view of cloud native deployments is useful as it sets up possible analogies to 5G.

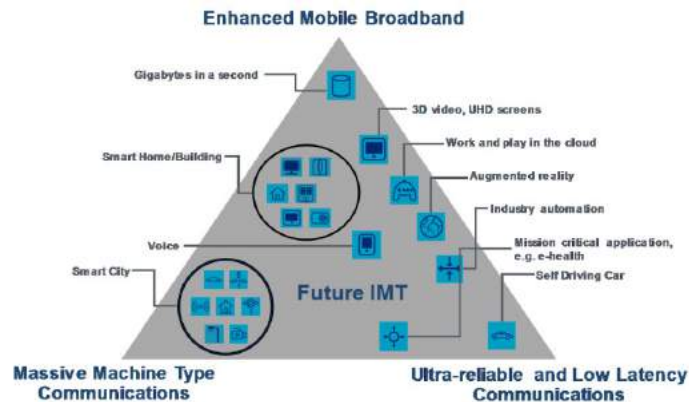


**Figure 4 - Typical Cloud Native Core Deployment**



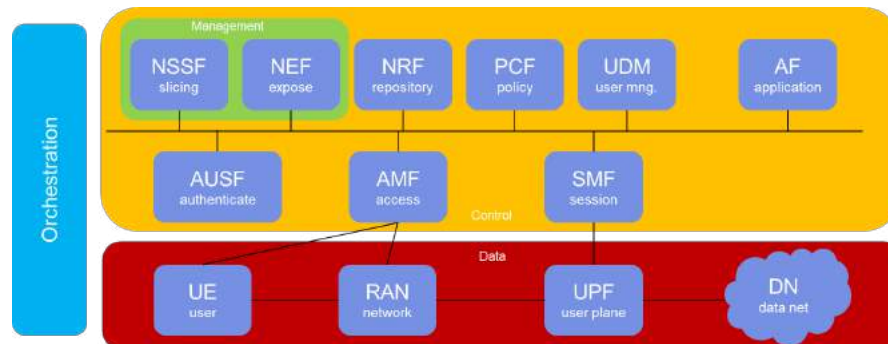
## 5. 5G and Network Slicing

The recent and very popular evolution to 5G in the mobile space provides us a good reference to learn from as we look at the convergence of multiple services and subservices. 5G is the mobile evolution from 4G LTE with drivers coming from increased bandwidth throughput, many more subscribed endpoints, and most importantly for us a widely varied set of usage scenarios, effectively creating a wealth of service types that have broadly varied expectations from the 5G network, see Figure 5. All these varied services are set to run simultaneously making the 5G network a largely different proposition than its predecessors.



**Figure 5 - 5G Usage Scenarios, (SANTO, 2017)**

The capability for 5G to carry various services types has been expressed not only in the radio access network, with many more cell sites and increased bandwidth but also in the way the 5G core has been architected. In Figure 6 we show the 5G core architecture definition according to the third generation partnership project (3GPP), (Mademann, 2017).



**Figure 6 - Cloud Native 5G Core**

A detailed review of the 5G NG architecture is beyond the scope of this paper but there are three notable evolutionary items making it different from its predecessor. They are listed below. In the next section we briefly describe each function.

- The 5G core is a cloud native, service based architecture, where architecture elements are defined as network functions with standard APIs such that they can be called on by any other network function with permission to do so (Felix, 2018). The standardized interfaces allow for multiple vendors to participate in the core, according to their capability in these discrete functions.

- 5G has separated the user plane and the control plane, allowing for top of the line hardware and software technologies and a flexibility of deployment models. This is referred to as Control User Plane Separation (CUPS). In this discussion we refer to the user plane as the data plane, see Figure 6.
- Most of the functions in the 5G architecture are carry over from LTE with some variations and containerization, however, some of them are new and specific to 5G. The new functions, we list under the management plane in Figure 6 are The Network Slice Selection function (NSSF), the Network Exposure Function (NEF). These functions were created specifically with the envisioned multiservice capability of 5G (variety of device types), and the dynamic nature of these services over time. Note that the calling out of NSSF/NEF as a management plane in Figure 6 is an exercise done for this paper to highlight the nature of these functions. The formal 5G architecture keeps them as just part of the control plane.

## 5.1. 5G Functions and Descriptions

Below is a brief description of the 5G core functions.


- UE: User Equipment. Any other device with mobile connectivity, such as smart phones.
- RAN: Radio Access Network. This is the network that connects user equipment to other parts of a mobile network via a radio connection.
- UPF: User Plane Function. Features to support packet routing and forwarding, interconnection to other data networks, and policy enforcement. Similar to the packet gateway in 4G LTE.
- DN: Data Network. Broader service provider network, “the internet.”
- AUSF: Authentication Server Function. Authenticates UEs and stores authentication keys.
- AMF: Access Management Function. manages user equipment registration, authentication, identification
- SMF: Session Management Function. Establishes and manages UE sessions, static and in movement, allocate IP addressing, informs on quality of service.
- PCF: Policy Control Function. Provides policy rules to control plane functions.
- UDM: Unified Data Management. Stores subscriber data and profiles.
- AF: Application Function. Connectivity point for management functions and control plane and data plane.
- NRF: Network Repository Function. Registration and discovery of network functions so that they can find each other.
- NSSF: Network Slice Selection Function. Has the task of selecting and directing network traffic to the use of particular network slice? Its assignments are determined by allowed usage per a network slice library found in the network slice selection assistance information (NSSAI). It also determines the access and mobility function (AMF) settings applicable to a user entity.
- NEF: Network Exposure Function. A mechanism that securely exposes state information of the 5G core, which includes capabilities and events, packet flow descriptors, and translation services for the flow of internal external information.

Of greatest interest in this discussion are the NSSF and the NEF, as together they allow for a dynamic network slicing mechanism over time.

## 5.2. What is Network Slicing

Network slicing is formally a method to run multiple end-to-end logical networks on a common set of resources. Network slicing in the most basic sense is not new. We have to date used to layer 1 slices with

managed OTN, or layer 2 slices with managed Ethernet Virtual LAN Networks (VLAN), or layer 3 with managed IP virtual private networks (VPN). This is certainly a part of network slicing, but in the more general case the slicing logic also happens end to end, which includes the partitioning of not only of packet network resources but also software control and management resources. Figure 7 depicts a specific set of resources being partitioned for the whole network. This end to end network slice would be available for a service to subscribe.



**Figure 7 - End to end network slicing representation**

There are particular service and operational expectations for 5G slices according to the 3GPP definition (3GPP, 2016). We summarize the expectations below to promote an intuitive understanding of how network slicing is leveraged in deployments. Ultimately as the concept of network slicing matures in the cable space there will need to be a well understood set of expectations, not exactly like the one listed below, but most likely similar with variations addressing the specific needs of MSO systems.

### **5.2.1. Service Requirements**

1. Slices are unique sets of network functions and configurations. The operator can create network slices from a multivendor environment of functions.
2. Network slicing is a dynamic exercise in an autonomous system applicable for diverse market scenarios.
3. A system can recognize a UE and its associated network slice.
4. A system allows the UE to subscribe to a specific network slice for service.

### **5.2.2. Operation Requirements**

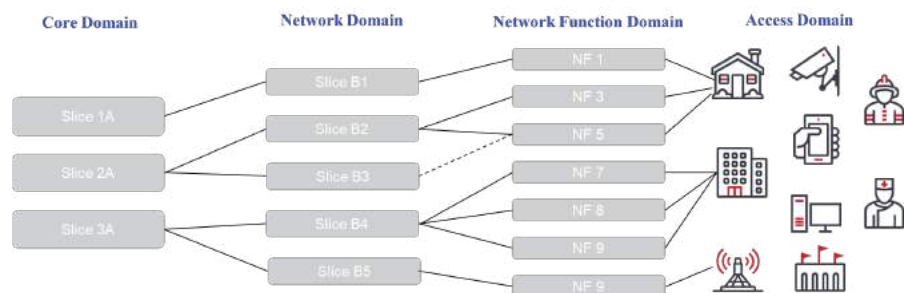
1. The operator can operate different network slices in parallel with isolation per slice.
2. Slices maintain the security profile expected of the service that uses it.
3. Network slices are isolated such that a cyber-attack would be confined to one slice.
4. Operators can authorize a third party to manage the network slicing environment, per suitable APIs.
5. The network slicing system is to scale in capacity with no impact to its service or other slices.
6. The system can accept changes to slices with minimal impact to subscriber services.

## **5.3. Service Convergence and Slicing**

A service is the instantiation of a set of features, policies, and configurations. A service is facilitated using a set of well-defined network functions and resources. In the context of network slicing we refer to a service as having a unique “blueprint”, and a blueprint can be paired to one or more slices.

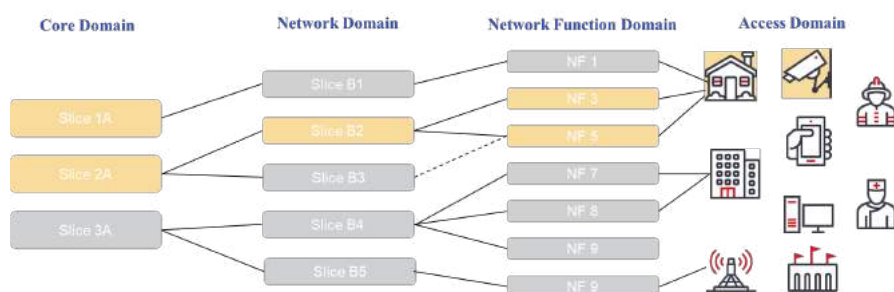
Service convergence then is the practice of multiple services (or subservices) coexisting in a well-defined manner on shared networking resources, (not unlike DOCSIS and Video do in legacy HFC). Service convergence is then simply the management of a collection of blueprints. Which implies, from previous sections, the management of parallel or interrelated slices.

For service convergence it is useful to take the birds eye view of a system, its multiple domains, slices, and their relationship. Figure 8 shows the components of a network slicing system for a typical service provider. There are several types of domains. The access domains with its multiple type of UEs which in the case of the MSO this would include the residential, enterprise and mobile access. We also have the network function domain, in practice this could include secure initialization, session bring up and tear down, encryption, etc. The network domain would be the resources for connectivity, in this domain there can be many slice types or various use cases available. This would be the case where you find distinctions for VLAN or VPNs, transmission methods for MPLS or Segment Routing (SR), each profile having its own slice. The core domain is a set of slices that give each service a unique profile southbound to the UE and interpret that service northbound to the operational and billing systems.



**Figure 8 - Domain and Slices Relationship**

What is important is not just the existence of slices but their inter-relationship in the system. We note that each service, or subservice can use a collection of network functions and subscribe to one or more network slices in a domain thus creating a specific blueprint. Figure 9, for example, shows a possible service blueprint in the highlighted mustard progression. Consider a security camera in the residential network. This camera needs resources from the network functions including the need for a timing service and local encryption. From the network domains it needs a VLAN or VPN tunnel, and from the core it needs a unique session and caching memory. This unique collection of resources is the blueprint for the security camera service.



**Figure 9 - Service Footprint**

While the functions, slice and slice pairings in Figure 9 are unique they are also dynamic in nature, as the collection of needed slices and network functions can change over the lifecycle of the service. Thus blueprints create scenarios where the relationship across domains can be collaborative, interdependent, and pre-configured to particular scenarios.


The configuration of unique scenarios, as the example given in Figure 9, along with the end to end network slice management necessary over its lifecycle is what is referred to network as a service,

commonly referred to as NaaS, where the creation of particular revenue streams is now possible because there is an infrastructure to create and manage different blueprints. This can certainly be a usable capability for the MSO space moving forward. Consider the blurring lines between home and work because of the COVID-19 pandemic. As we considered the earlier example there is a case to differentiate, though the whole network, work traffic at home versus entertainment traffic at home—effectively two different blueprints.

## 6. MSO and 5G coexistence

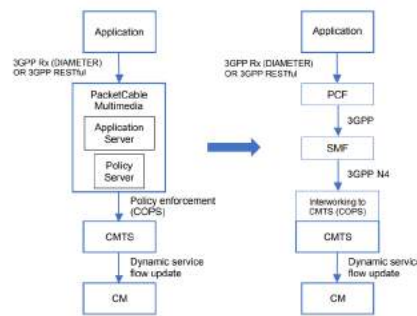
In Figure 1, of section 3.1 we supposed the reuse of redundant functions in service cores, thus creating end-to-end convergence. But this goal is aspirational, and to think about an interim path is necessary. We refer to the work done to support the momentum for integration of small cells onto the cable landscape, with DOCSIS backhaul (CableLabs, 2020). Per definition there is a generic satellite function to the 5G core that allows for a wireline system, to present itself as a user entity to the 5G core (3GPP, 2018). This function is called the Wireline Access Gateway Function, (W-GAF) and has standard interfaces, transmitting data plane traffic to the user plane function and control traffic to the 5G core control plane. This in effect allows for any non-5G system to have an embedded RAN and subtended devices, with the expectation that the wireline core has the capability to decipher separate data and control plane traffic.

In the case of cable, as shown in Figure 10, the W-GAF facilitates having small cells within the CMTS/DOCSIS system and creates an environment for UE's, behind the CMTS, to be treated as native participants of the 5G core. This is made possible in conjunction with the work done for DOCSIS MAC timing and latency functions (Cablelabs Timing, 2020) (Cablelabs Xhaul, 2020)



**Figure 10 - CMTS / 5G cores coexistence**

The working model for the coexistence of the CMTS and 5G, via a standardized gateway, gives a glimpse on the possibility to reuse the existing NSSF and NEF as a generalized slicing solution for other cores. The work at 3GPP championed by CableLabs, reported in the document “5G Wireless Wireline Converged Core Architecture Technical Report” (CableLabs, 2020) shows that there is already an established data model for an HFC system with a RAN that can subscribe to, be authenticated by, and admitted into the 5G core. Further there is a proposed method to extend not only subscription to the core but extend the QoS mechanism from the 5G core, available through the wireline gateway to the CMTS and its constituents. This is somewhat straight forward as the RAN components in the HFC have native wireless tendencies towards the control and data plane structure of the 5GC.



**Figure 11 - General QoS mapping CMTS and 5G core (CableLabs, 2020)**

Outstanding however is how to incorporate the non-wireless traffic into the global view for network slicing. What would be useful is a global QoS mapping to all CMTS traffic, as expressed in the DOCSIS 3.1 MAC Upper Layer Protocols Interface (MULPIv3.1) and Packet Cable Multimedia (PCMM), creating a complete map for all services, see Figure 11. The extension of QoS for native DOCSIS traffic can then be interpreted by 5G core Session Management Function (SMF) via a Common Open Policy Service (COPS) (CableLabs, 2020). With a global QoS structure in place the creation of MSO non-mobile specific slice types can be defined. The management of both HFC and 5G slices can then be done by the already existing SFFM.

It is worth mentioning that now QoS tagging is generally not a practice for DOCSIS packets in general and not over the new digital access fiber infrastructure. This is because it relies on the time map mechanism between CMTS and modems. There is a QoS mechanism that is embedded into the Xhaul specification, which will be necessary when transmitting small cells over DOCSIS and could prove useful if there is a general application of QoS for all residential signals would be necessary. An extension of QoS tagging to the Ethernet / IP network could prove useful or necessary as the MAC layer evolves in positioning, per the efforts of the FMA. In these cases, the network can participate in more intricate slices along with other elements of the end-to-end network.

An extension of this principle can be applied to the enterprise vertical, adopting enterprise bandwidth profiles to the global QoS mapping, starting at the first network aggregation point. The network slices for the enterprise would in principle use the structure laid out by native OAM signaling.

## 7. Slicing methods

### 7.1. Soft Slicing

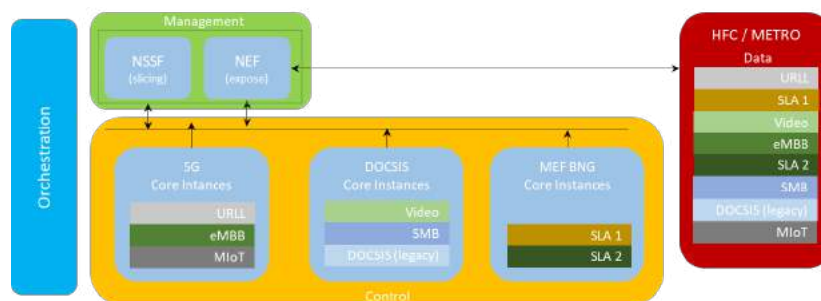
Soft slicing refers to the practice of sharing resources with system enabled flexibility. In soft slicing the slices are statistically separate and transmit by logically multiplexing the data plane over some physical channel (IETF, 2018). Soft slicing makes use of the layer 2 Ethernet and layer 3 Internet Protocols we have grown accustomed to, dynamic MPLS, Segment Routing, other tunneling mechanisms, and even RAN frequency sharing.

Figure 12 shows a soft slicing example for convergence, with 5G, DOCSIS and MEF Enterprise services represented. Note that within each core there is a QoS priority structure, native to each one and the cores may or may not be sharing functional resources on the control plane. On the data plane however we see that there is an expression of an overall QoS per the management of the NSSF and the dynamic



management of the NEF. This exercise implies that the networking infrastructure is programmable with stateful telemetry to the management systems.

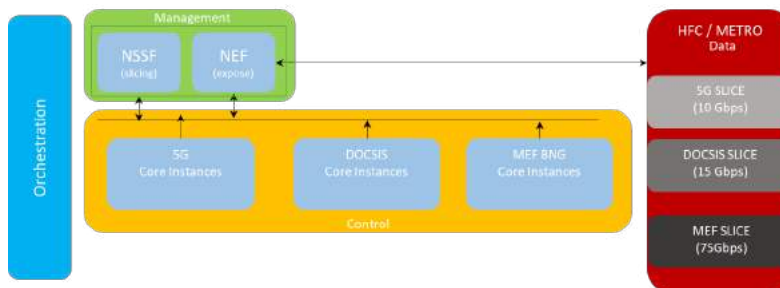
Soft slicing is useful when the operator is trying to maximize the return on investment on networking equipment or trying to minimize carbon and physical footprint of the data plane due to other forces. Note that the prioritization in Figure 12 necessarily implies a policy mechanism by the operator that sets the priority structure for the signaling and while technically possible it can have its hurdles per historical precedents discussed in Section 3.



**Figure 12 - Soft Slicing Example for Convergence**

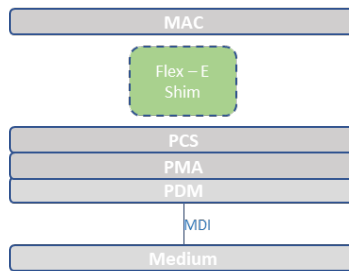
## 7.2. Hard Slicing

Hard slicing refers to the provision of resources in such a way that they are dedicated to a blueprint instantiation. This refers to a slice being assigned a particular lambda or ONT muxponder or Flex Ethernet (FlexE) channel.



**Figure 13 - Hard Slicing Example for Convergence**

Figure 13 shows an example of convergence through hard slicing, using the mechanism of FlexE, as defined in the specifications from the Optical Internet Forum (OIF) and International Telecommunications Union (ITU): OIF-FLEXE-02.01A and ITU-T G.mtn, respectively. FlexE is basically a time domain multiplexing of Ethernet signals where, as part of its capabilities, signals of lower bandwidths can aggregate to a higher bandwidth signal and vice versa, this channelization effect can be maintained throughout the Ethernet deployment, with 5 Gbps of granularity. Formally FlexE adds a flexible shim layer between the MAC and the physical coding sublayer, see Figure 14. Note that the implementation of FlexE is in a large part made possible by the arrival and availability of ZR type coherent optics in the range of 100-400 Gbps, which allow for robust access backhaul solutions.

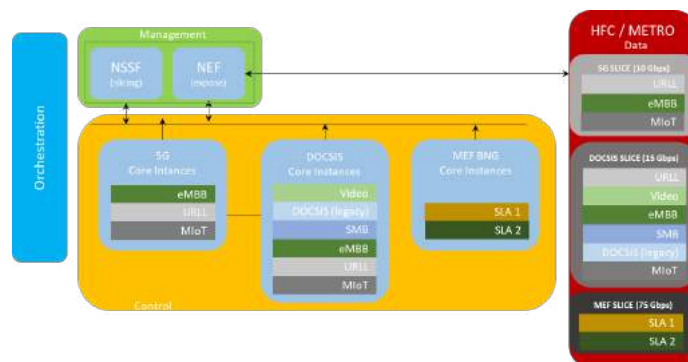


**Figure 14 - Flexible Ethernet Operational Definition**

In the Figure 13 example there is no system wide knowledge of the prioritization of the subservices within the 5G, DOCSIS or Enterprise services. Instead there is a bandwidth expectation for each service and a FlexE channelization assignment and an understanding from the NSSF as to the assigning of bandwidth per service. The bandwidth assignments themselves are flexible over the lifecycle of the service, however. Hard slicing is useful in cases where resource sharing is needed but hard boundaries between them are necessary. In the MSO case, this could prove a worthwhile steppingstone as the industry moves away from siloed operations per service.

### 7.3. Hard And Soft Slicing

The combination of hard and soft slicing for the MSO operator could be an implementable middle ground to optimize resources and create the right sort of boundary structure in the first implementation of convergence. As shown in Figure 14 soft slicing is limited to the purview of subservices within each core and in this manner there is no overreaching policy that pins a subservice on one core versus a subservice on a different core. Simultaneously an implementation of FlexE per core is introduced such that each core has a dedicated bandwidth to work with, allowing no confusion on usage limitation per QoS assignments of its subservices. In this case the NSSF is executing on blueprints that include several protocol layers, and as is expected the system is dynamic with the capabilities to evolve hard and soft slicing assignments over time. For the MSOs this can be a long term or near-term solution for convergence.



**Figure 15 - Hard and Soft Slicing Example**



## 8. Network Slice Lifecycle

It is useful to understand the lifecycle of a network slice. Note that the existence of the network exposure function is key here as by its gathering and maintaining of stateful system information the capacity of slices can be anticipated and thus react accordingly. Below we propose the 5 steps necessary for a lifecycle. These steps help the operator to appreciate the effort necessary in dealing with a network sliced system and creating an expectation for native or third-party enablement of these services.

- Creation
  - In this step a machine-readable blueprint definition is created. This is created by an operator or a third-party planner. The blueprint includes component resources, enabling features, workflow assignments and lifecycle expectation.
- Instantiation
  - At this step resource inventory and availability are counted. It is where function discovery of native or third party VNFs take place, and the orchestration of the system pushes for creation of slices.
- Scaling
  - This includes monitoring network throughput and trigger events. Allows for the extension or reduction of slice capabilities per need. All this is done with zero touch, it is an automated process.
- Isolation
  - This step manages resource impact of scaling on neighbor slices, including items like traffic, bandwidth among other processes. It insures parallel usage.
- Maintenance
  - This step facilitates in instantiation or tear down of slices, redirecting traffic to alternate slices as necessary with minimal or no interruption, and conducts proactive testing.

## 9. Industry Recommendations

We list several helpful steps that are meant to help the industry facilitate convergence. The list below is non-exhaustive but does provide the industry with goals worth considering.

- Move towards industry-based open interfaces for cloud native core functions.
  - In the minimum open interfaces for network slicing like functions and network exposure like functions.
- Practice QoS and involve the packet network.
  - Even though QoS capabilities are inherent in DOCSIS for instance, they are not used all the time, nor are they generally reflected in the digital access network.
- Create MSO specific list of slice expectations and requirements.
  - This was an initial and necessary step in the creating of network slicing for the mobile industry. This is an activity CableLabs could lead for example.
- Implement hard or hard and soft slicing as a beginning step.
  - The cable industry must work with, not against, the reality of heavily siloed services, and move from there. Hard slicing is a useful tool in this direction.
- Move towards full functional convergence at the core.
  - This is a long-term goal, but it begins with definitions. The work done in the FMA of standard interfaces for certain MAC functions is a good example of how to start. Ultimately, legacy cable functions are competing with a system in 5G that is written for cloud and with standardized interfaces.

## 10. Conclusion

The topic of network slicing for convergence of services is novel and necessary. In this paper we have proposed a framework for service convergence using network slicing. We have reviewed the network slicing mechanisms for 5G and pointed out possible analogies that aid in developing slicing for MSO systems containing residential, business, and mobile services. We have covered the concepts of network slicing functions which organize and partition available network resources. We have described hard and soft slicing mechanisms and the necessary steps to maintain end-to-end slice visibility and usability over their lifecycle.

Lastly, we acknowledge the fruitful discussions we've had developing the content of this paper with Raghu Ranganathan and Darren McKinney from Ciena, along with Bernard McKibben from CableLabs.

## Abbreviations

|        |                                                 |
|--------|-------------------------------------------------|
| 3GPP   | 3rd Generation Partnership Project              |
| 5G     | 5th Generation Mobile Network                   |
| AF     | Application Function                            |
| AMF    | Access Management Function                      |
| API    | Application Programming Interface               |
| AUSF   | Authentication Server Function                  |
| BNG    | Broadband Network Gateway                       |
| CCAP   | Converged Cable Access Platform                 |
| CMTS   | Cable Modem Termination System                  |
| CPE    | Customer Premise Equipment                      |
| CUPS   | Control User Plane Separation                   |
| DAA    | Distributed Access Architecture                 |
| DN     | Distribution Network                            |
| DOCSIS | Data Over Cable Service Interface Specification |
| eMBB   | Enhanced Mobile Broadband                       |
| FlexE  | Flex Ethernet                                   |
| FMA    | Flexible MAC Architecture                       |
| GbE    | Gigabit Ethernet                                |
| Gbps   | Gigabits per second                             |
| HFC    | Hybrid Fiber Coaxial network                    |
| IP     | Internet Protocol                               |
| ISBE   | International Society of Broadband Experts      |
| MAC    | Media Access Control                            |
| MEF    | Metro Ethernet Forum                            |
| MIoT   | Massive Internet of Things                      |
| MPEG   | Moving Picture Experts Group                    |

|       |                                            |
|-------|--------------------------------------------|
| MPLS  | Multiprotocol Label Switching              |
| MSO   | Multiple-System Operator                   |
| NaaS  | Network As A Service                       |
| NEF   | Network Exposure Function                  |
| NRF   | Network Repository Function                |
| NSSF  | Network Slice Selection Function           |
| OLT   | Optical Line Terminal                      |
| OPEX  | Operational Expenses                       |
| OTN   | Optical Transport Network                  |
| PCF   | Policy Control Function                    |
| PON   | Passive Optical Network                    |
| QAM   | Quadrature Amplitude Modulation            |
| QoS   | Quality of Service                         |
| RAN   | Radio Access Network                       |
| RF    | Radio Frequency                            |
| RMD   | Remote MAC-PHY Device                      |
| RPD   | Remote PHY Device                          |
| SCTE  | Society of Cable Television Engineers      |
| SFM   | Session Management Function                |
| SLA   | Service Level Agreement                    |
| SMB   | Small Medium Business                      |
| SR    | Segment Routing                            |
| TDM   | Time Domain Multiplexing                   |
| UDM   | Unified Data Management                    |
| UE    | User Equipment                             |
| UPF   | User Plane Function                        |
| URLLC | Ultra Reliable Low Latency Communication   |
| VNF   | Virtual Network Function                   |
| VSG   | Vertical Systems Group                     |
| Wi-Fi | Wireless Fidelity based on the IEEE 802.11 |

## Bibliography

3GPP. (2016). *3rd Generation Partnership Project TR 22.891 V14.2.0*. Valbonne: 3GPP Organizational Partners.

3GPP. (2018). *Wireless and wireline convergence access support for the 5G System*. Valbonne: 3rd Generation Partnership Project.

Andreoli-Fang, J. (2019, September 10). *Enabling 5G with 10G Low Latency Xhaul (LLX) Over DOCSIS® Technology*. Retrieved from Informed Blog by Cablelabs:  
<https://www.cablelabs.com/enabling-5g-10g-low-latency-xhaul-llx-docsis-technology>

- BAUMGARTNER, J. (2019, April 16). *Comcast, Charter MVNO Deals Are Bad for Everyone – Analyst*. Retrieved from Light Reading: JEFF BAUMGARTNER
- BTR . (2019, December 24). *Vecima demos FMA API interoperability for CableLabs*. Retrieved from Broadband Technology Report: <https://www.broadbandtechreport.com/docsis/article/14074102/vecima-demos-fma-api-interoperability-for-cablelabs>
- CableLabs. (2020). *5G Wireless Wireline Converged Core Architecture Technical Report*. Louisville: Cable Television Laboratories Inc.
- Cablelabs Timing. (2020). *Synchronization Techniques for DOCSIS® Technology*. Louisville: Cable Television Laboratories, Inc.
- Cablelabs Xhaul. (2020). *Low Latency Mobile Xhaul over DOCSIS Technology*. Louisville: Cable Television Laboratories, Inc.
- Chamberlain, J. (2018, January 16). *Necessities for Network Convergence in 2018 and Beyond (Part 1)*. Retrieved from Commscope: <https://www.commscope.com/blog/2018/necessities-for-network-convergence-in-2018-and-beyond-part-1/>
- Charter . (2020, July 31). *Charter Announces Second Quarter 2020 Results*. Retrieved from Charter Communications News: <https://ir.charter.com/static-files/8402d27e-e891-41ce-ba11-b8fd55f79709>
- Comcast . (2020, July 30). *Comcast Reports 2nd Quarter 2020 Results*. Retrieved from Investor News Details : <https://www.cmcsa.com/news-releases/news-release-details/comcast-reports-2nd-quarter-2020-results>
- Felix, E. (2018, October 20). *5G Service-Based Architecture (SBA)*. Retrieved from Medium: <https://medium.com/5g-nr/5g-service-based-architecture-sba-47900b0ded0a>
- Hodges, J. (2019, June 4). *The Rise of Network-as-a-Service*. Retrieved from Light Reading Cloud Services: <https://www.lightreading.com/services/cloud-services/the-rise-of-network-as-a-service/a/d-id/752185>
- IETF. (2018, January 4). *Network Working Group*. Retrieved from Network Slicing Architecture, draft-geng-netslices-architecture-02: <https://tools.ietf.org/id/draft-geng-netslices-architecture-02.html#rfc.section.4.3>
- Mademann, F. (2017, December 21). *System architecture milestone of 5G Phase 1 is achieved*. Retrieved from 3GPP A Global Initiative: [https://www.3gpp.org/news-events/1930-sys\\_architecture](https://www.3gpp.org/news-events/1930-sys_architecture)
- SANTO, B. (2017, August 25). *The 5 best 5G use cases*. Retrieved from EDN: <https://www.edn.com/the-5-best-5g-use-cases/>
- Vertical Systems Group. (2020, April 22). *2019 U.S. Carrier Managed SD-WAN LEADERBOARD*. Retrieved from Vertical Systems Group: <https://www.verticalsystems.com/2020/04/21/2019-us-sd-wan-leaderboard/>
- Villarruel, F. (2014). Plasticity of the New HFC Network Engineering for Remote-PHY and FTTP. *SCTE Cable Tech Expo 14* (pp. 1-23). Denver: SCTE.

Villarruel, F. (2015). Virtual PON Network A Practical Guide For The Network Planner. *Cable-Tec Expo 15* (pp. 1-22). New Orleans: SCTE.

Villarruel, F. (2018). Capacity and Technology Considerations in DAA. *SCTE ISBE NCTA CABLELABS 2018 Fall Technical Forum* (pp. 1-23). Atlanta: SCTE\*ISBE.

# **Satellite to Fiber Broadcast Execution With SCTE 35 and 224**

## **Implementing Linear Rights and Alternate Content for Network and Affiliate Feeds**

A Technical Paper prepared for SCTE•ISBE by

### **Stuart Kurkowski, PhD**

Distinguished Engineer, Principal Architect  
Comcast Technology Solutions  
Dry Creek Facility  
303-503-2680  
Stuart\_Kurkowski@comcast.com

### **Neill Kipp**

C&SP Architect Lead  
Comcast Technology Solutions  
1899 Wynkoop Street, Suite 500  
720-530-6917  
Neill\_Kipp@comcast.com

# Table of Contents

| Title                                      | Page Number |
|--------------------------------------------|-------------|
| 1. Introduction.....                       | 4           |
| 2. Standards Background.....               | 5           |
| 2.1. ANSI/SCTE 224 .....                   | 5           |
| 2.1.1. Media .....                         | 6           |
| 2.1.2. MediaPoints .....                   | 6           |
| 2.1.1. Policy.....                         | 6           |
| 2.1.1. ViewingPolicy .....                 | 6           |
| 2.1.1. Audience .....                      | 7           |
| 2.2. ANSI/SCTE 35 .....                    | 7           |
| 2.2.1. ANSI/SCTE-35 Serialization .....    | 7           |
| 2.2.2. XML .....                           | 9           |
| 2.3. ANSI/SCTE 250 .....                   | 9           |
| 2.3.1. Message Transport .....             | 10          |
| 2.3.2. Message Namespace .....             | 10          |
| 2.3.3. Message Time.....                   | 10          |
| 2.3.4. ESAM Interface .....                | 11          |
| 2.3.5. In-band ESAM Message Exchange ..... | 11          |
| 2.4. Use Cases .....                       | 11          |
| 2.4.1. Decision Request .....              | 11          |
| 2.4.2. Decision Response .....             | 12          |
| 2.4.3. SAS Response Acknowledgement .....  | 13          |
| 3. Distribution Workflow .....             | 13          |
| 3.1. Transition with Automation .....      | 14          |
| 3.1.1. Signaling .....                     | 14          |
| 3.1.2. Additional Instruction.....         | 15          |
| 3.2. Transition without Automation.....    | 15          |
| 3.2.1. Signal Insertion .....              | 15          |
| 3.2.1. Additional Instruction.....         | 17          |
| 4. Solution .....                          | 17          |
| 5. Additional thoughts.....                | 17          |
| 5.1. OTT Complexity .....                  | 17          |
| 5.2. EPG Usage .....                       | 18          |
| 6. Conclusion.....                         | 18          |
| Abbreviations .....                        | 20          |
| Bibliography & References.....             | 20          |
| Appendix A.....                            | 22          |

## List of Figures

| Title                                                      | Page Number |
|------------------------------------------------------------|-------------|
| Figure 1 - High Level Architecture .....                   | 5           |
| Figure 2 - Example ANSI/SCTE 224 Policy Object.....        | 6           |
| Figure 3 - Example ANSI/SCTE 224 ViewingPolicy Object..... | 7           |
| Figure 4 - Example ANSI/SCTE 224 Audience Object .....     | 7           |
| Figure 5 - Example ANSI/SCTE 35 Hexadecimal Tag .....      | 8           |

|                                                                             |    |
|-----------------------------------------------------------------------------|----|
| Figure 6 - Example ANSI/SCTE 35 Base64 Tag .....                            | 9  |
| Figure 7 - Example ANSI/SCTE 35 M3U8 Base64 Tag .....                       | 9  |
| Figure 8 - Example ANSI/SCTE 35 XML Serialization .....                     | 9  |
| Figure 9 - Traditional ESAM Exchange.....                                   | 10 |
| Figure 10 - Example ESAM SignalProcessingEvent .....                        | 12 |
| Figure 11 - Example ESAM SignalProcessingNotification Noop .....            | 12 |
| Figure 12 - Example ESAM SignalProcessingNotification for a Switch.....     | 13 |
| Figure 13 - Example ESAM SignalProcessingNotification Acknowledgement ..... | 13 |
| Figure 14 – Detailed Workflow .....                                         | 14 |
| Figure 15 - Example MediaPoint MatchSignal.....                             | 15 |
| Figure 16 - Workflow for Signal Insertion .....                             | 16 |

## List of Tables

| <b>Title</b>                                    | <b>Page Number</b> |
|-------------------------------------------------|--------------------|
| Table 1 - Example ANSI/SCTE 35 UPID Types ..... | 8                  |
| Table 2 - Example ANSI/SCTE 35 Namespaces.....  | 10                 |



# 1. Introduction

The world of broadcast video delivery is evolving rapidly, especially with the proceeding by the Federal Communications Commission (FCC) to clear the lower 280 MHz of the C-band spectrum, to pave the way for its use by 5G services (FCC Report and Order). Satellite companies including Intelsat and SES use the C-band spectrum to serve TV broadcasters and cable network operators with video feeds. The FCC ordered these companies to move their operations to the upper 200 MHz of C-band in order to prevent interference from mobile services (Reardon 2020).

As a result of the uncertainty of C-band going forward and the continual demands for higher video quality, broadcasters and content providers are considering transitioning more fully from satellite to fiber, and/or Internet-based delivery mechanisms. Linear channel delivery over terrestrial optical fiber offers lower cost, lower latency, lower bandwidth constraints, and higher quality than a corresponding delivery over satellite. With these benefits and the FCC mandates, operators have some compelling reasons to make the transition from satellite to fiber.

Any transition to terrestrial video delivery over fiber must occur with minimal impact to current customer workflows. The risks of transition certainly include disruption of primary service, but also of signaling for advertising, local contributions, and alternate programming, including blackouts. The benefits include additional over-the-top (OTT) distribution. Certain hybrid use cases increase disruption risk. For example, national feeds arriving by satellite can signal a local contribution arriving over fiber, or vice versa.

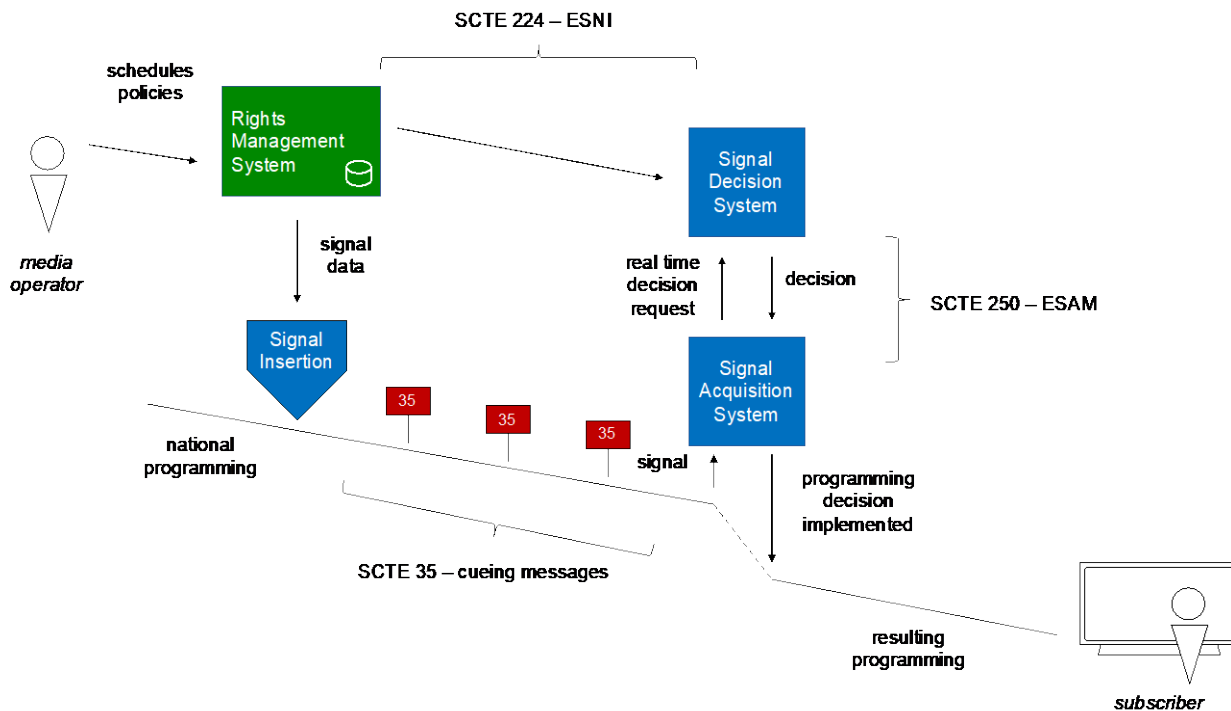
The SCTE Digital Video Subcommittee (DVS) maintains a family of standards that, when implemented together, help support a seamless transition from satellite to fiber.

Regardless of transmission path, linear streams carry in-band signaling following ANSI/SCTE 35, the “Digital Program Insertion Cueing Message for Cable” standard. ANSI/SCTE-35 signals convey the start and end of advertising, programming boundaries, and license boundaries.

The “Event Scheduling and Notification Interface” (ESNI; ANSI/SCTE-224) is an out-of-band signaling standard that allows the enforcement of viewing restrictions, such as web embargos, startover, lookback, and alternate programming. ANSI/SCTE-224 provides the mechanism to achieve the video markup and execution for all aspects of the workflow, including national and local contributions, ad insertion, as well as execution on the playout side for local broadcasts and OTT content.

And, the “Real-time Event Signaling and Management API” (ESAM; ANSI/SCTE-250) provides a software standard for converting in-band ANSI/SCTE-35 or out-of-band ANSI/SCTE-224 signals into real-time requests of a signal decision system. ANSI/SCTE-250 can be used to implement advertising placements, license constraints, and local programming.

A linear signal processing system that utilizes the combination of ANSI/SCTE-35, 224, and 250 requires a responsive and scalable architecture that can be implemented using a combination of on-premise systems and cloud-based services (Figure 1). Implementing this linear system enables the seamless transition of fully featured linear channels from satellite to the Internet.



**Figure 1 - High Level Architecture**

## 2. Standards Background

The interplay between content rights and multi-platform distribution has become increasingly complex. Every component in a delivery workflow must weigh simplicity and deployment agility against quality of experience. Today, configuration data is manually entered into spreadsheets and application programming interfaces (APIs) that are distributed via multiple channels. Programmers are reliant on operators to acknowledge that data has been received, and then operators must ingest data from multiple sources and formats before it can be assembled and delivered. By utilizing standards created by the SCTE DVS and specifically Work Group 5, Digital Program Insertion, all parts of the workflow can work together at a machine-to-machine level. In the following sections we detail these various standards that enable this machine-to-machine workflow.

### 2.1. ANSI/SCTE 224

ANSI/SCTE 224 ESNI provides a robust framework of Extensible Markup Language (XML) messages. It details descriptions of audience characteristics, and viewing policies associated with each audience, it also allows for channels (Media) and individual events (MediaPoints) to describe start and end times (MatchTime) or in-band signaling (MatchSignal) information. With ANSI/SCTE 224, operators can implement bindings from events and audiences to specific programming signals; this will signal downstream systems the desired programming changes and exactly when to implement them. Although the framework is robust, operators must assemble meticulous programming details into a standard format, control metadata delivery targets, maintain metadata visibility across a diverse delivery ecosystem, and reliably communicating the resulting instructions to every relevant distribution point.

ANSI/SCTE 224 allows for rights convenience by setting up five different message types to carry the information: `Media`, `MediaPoints`, `Policy`, `ViewingPolicy`, and `Audience`. The following sections detail each of these message types and provide a brief description of each.

### **2.1.1. *Media***

The `Media` object represents a linear television channel and its schedule of events. The `Media` is the container that contains the `MediaPoints` used to describe individual events that are scheduled to occur on that channel.

### **2.1.2. *MediaPoints***

`MediaPoints` are individual events that occur within a channel or a show. These events are traditionally program start, ad start, ad stop, and program end. The `MediaPoint` is matched against an event in the video either by time-based triggers (e.g., program start at 2 PM) or matching against the SCTE 35 in-band signal described below. Using the in-band signal has the advantage of allowing the actions to be taken on a frame-accurate basis, making for a better user experience. When a `MediaPoint` is matched to an in-video event, it is “triggered,” which means that one or multiple `Policy` messages can be applied to the state of the channel. The `MediaPoint` can also contain additional metadata such as alternate identifiers and electronic program guide (EPG) information. This enables live to video on demand capture or accurate EPG usage with the out-of-band ANSI/SCTE 224 data. See Appendix A for a complete example of a ANSI/SCTE 224 `Media` object.

### **2.1.1. *Policy***

The `Policy` message is a container for `ViewingPolicy` that can be associated with an event (Figure 2). `Policy` elements are added and removed to control channel playout. For example, a `Policy` about trick mode restrictions can be added to the channel state and a second `Policy` that sends an audience to an alternate channel can be added as well. It is also at the `Policy` level that items are removed from the state, so in the previous example, if the `Policy` to send the audience to an alternate channel is removed from the state, then that audience will return to watching the main channel.

```
<Policy xmlns="http://www.scte.org/schemas/224"
 id="xyz/policy/any/slate"
 lastUpdated="2016-11-18T15:00:00.000Z">
 <ViewingPolicy xmlns="http://www.scte.org/schemas/224"
 xmlns:xlink="http://www.w3.org/1999/xlink"
 xlink:href="xyz/viewingpolicy/any/slate">
 </ViewingPolicy>
</Policy>
```

**Figure 2 - Example ANSI/SCTE 224 Policy Object**

### **2.1.1. *ViewingPolicy***

`ViewingPolicy` elements are key messages, because they bring together the action for a triggered event and the appropriate audience. So, for example, the `ViewingPolicy` for one audience might be to message telling the view the programming is unavailable, while a different audience might be sent to an alternate feed (Figure 3). Defined sets of actions cover content directives, trick mode restrictions, dynamic ad insertion, capture controls for recording, and signal insertion and deletion.

```

<ViewingPolicy xmlns="http://www.scte.org/schemas/224"
 id="xyz/viewingpolicy/any/slate"
 lastUpdated="2016-11-18T15:00:00.000Z">
 <Audience xmlns="http://www.scte.org/schemas/224"
 xmlns:xlink="http://www.w3.org/1999/xlink"
 xlink:href="xyz/audience/location/any">
 </Audience>
 <Content xmlns="urn:scte:224:action">urn:scte:224:action:slate</Content>
</ViewingPolicy>

```

**Figure 3 - Example ANSI/SCTE 224 ViewingPolicy Object**

### **2.1.1. Audience**

Audience is individual messages used to characterize the groups of viewers impacted by a ViewingPolicy (Figure 4). Numerous standard audience characterization elements include ZIP codes, latitude and longitude, device types, and viewer status.

```

<Audience xmlns="http://www.scte.org/schemas/224"
 id="xyz/audience/location/any"
 description="No Travel Restrictions"
 lastUpdated="2016-10-22T16:00:00.000Z"
 match="NONE">
 <Zip xmlns="urn:scte:224:audience">00000</Zip>
</Audience>

```

**Figure 4 - Example ANSI/SCTE 224 Audience Object**

Combining these five message types into sets of machine-to-machine executable instructions is the key to utilizing ANSI/SCTE 224 to manage the execution of content, at scale, across multiple distribution partners.

## **2.2. ANSI/SCTE 35**

ANSI/SCTE-35 in-band signals appear in MPEG-2 video transport streams and narrate events as they play out. ANSI/SCTE-35 signals primarily mark when programs start or stop and when ad insertion should start and stop. Signals can also indicate emergency action messages. Each signal has an eight-bit segmentation\_type\_id that tells the receiver what type of event is being signaled (0x00 - 0xFF).

For example, if a baseball game started at 1 PM and ended at 3:45 PM, inside the stream would be a 0x10 signal at 1 PM for the program start and a 0x11 signal at 3:45 PM for the program end (and then probably a 0x10 signal at the same time for the next program start). There would be many 0x34/0x35 pairs that indicate when advertising placement opportunities appear. Program boundary signals 0x10 and 0x11 contain program identifiers, also known as unique program identifiers (UPID), that bind to ANSI/SCTE-224 rules based on programs.

In other words, programmers notify downstream devices by putting ANSI/SCTE-35 messages in a stream to signal ANSI/SCTE-224 channel events.

### **2.2.1. ANSI/SCTE-35 Serialization**

ANSI/SCTE-35 has two basic serialization forms. The most basic form is a raw binary sequence. Table 1 shows an example. What's shown is just part of the SCTE-35 message, not the whole container.

**Table 1 - Example ANSI/SCTE 35 UPID Types**

| Syntax                          | Bits | Mnemonic | Encrypted |
|---------------------------------|------|----------|-----------|
| splice_info_section() {         |      |          |           |
| table_id                        | 8    | uimsbf   |           |
| section_syntax_indicator        | 1    | bslbf    |           |
| private_indicator               | 1    | bslbf    |           |
| reserved                        | 2    | bslbf    |           |
| section_length                  | 12   | uimsbf   |           |
| protocol_version                | 8    | uimsbf   |           |
| encrypted_packet                | 1    | bslbf    |           |
| encryption_algorithm            | 6    | uimsbf   |           |
| pts_adjustment                  | 33   | uimsbf   |           |
| cw_index                        | 8    | uimsbf   |           |
| tier                            | 12   | bslbf    |           |
| splice_command_length           | 12   | uimsbf   |           |
| splice_command_type             | 8    | uimsbf   | E         |
| if(splice_command_type == 0x00) |      |          |           |
| splice_null()                   |      |          | E         |
| if(splice_command_type == 0x04) |      |          |           |
| splice_schedule()               |      |          | E         |
| if(splice_command_type == 0x05) |      |          |           |
| splice_insert()                 |      |          | E         |
| if(splice_command_type == 0x06) |      |          |           |
| time_signal()                   |      |          | E         |
| if(splice_command_type == 0x07) |      |          |           |
| bandwidth_reservation()         |      |          | E         |
| if(splice_command_type == 0xff) |      |          |           |
| private_command()               |      |          | E         |
| descriptor_loop_length          | 16   | uimsbf   | E         |
| for(i=0; i<N1; i++)             |      |          |           |
| splice_descriptor()             |      |          | E         |
| for(i=0; i<N2; i++)             |      |          |           |
| alignment_stuffing              | 8    | bslbf    | E         |
| if(encrypted_packet)            |      |          |           |
| E_CRC_32                        | 32   | rpchof   | E         |
| CRC_32                          | 32   | rpchof   |           |
| }                               |      |          |           |

This signal can also be expressed in hexadecimal, as in a #EXT-X-DATERANGE tag in M3U8 (Figure 5).

```
#EXT-X-DATERANGE:ID="splice-6FFFFFF0",START-DATE="2014-03-05T11:15:00Z",PLANNED-
DURATION=59.993,SCTE35-
OUT=0xfc304c0000000000ffff00506fee36333f7003602194355454900882da97f80010a3030303839
3234353835100000021943554549007ef2567f80010a3030303833313935373411000058b3f887
```

**Figure 5 - Example ANSI/SCTE 35 Hexadecimal Tag**

It can also be expressed as Base64, as in ESAM messages, as described in section 2.3 (Figure 6).

```
<signaling:Binary
signalType="SCTE35">/DBMAAAAAAAAA//wBQb+42Mz9wA2Ah1DVUVJAigtqX+AAQowMDA4OTI0NTg1EAAAAh
1DVUVJAH7yVn+AAQowMDA4MzE5NTc0EQAAWLP4hw==</signaling:Binary>
```

**Figure 6 - Example ANSI/SCTE 35 Base64 Tag**

The M3U8 EXT-X-SCTE35 and EXT-OATCLS-SCTE35 tags also use a Base64 implementation (Figure 7).

```
#EXT-X-
SCTE35:TYPE=0x10,TIME=1450707450000,ELAPSED=0,ID="dAQpTUaQSjOti/JZTqECfQ==",CUE="/DBMA
AAAAAAAA//wBQb+42Mz9wA2Ah1DVUVJAigtqX+AAQowMDA4OTI0NTg1EAAAAh1DVUVJAH7yVn+AAQowMDA4MzE
5NTc0EQAAWLP4hw=="
...
#EXT-OATCLS-SCTE35:/DA1AAAAAAAAAAP/wFAXwAAACf+/+AM3Qkf4AFJlwAAEBAQAAGkXGSg==
```

**Figure 7 - Example ANSI/SCTE 35 M3U8 Base64 Tag**

### 2.2.2. XML

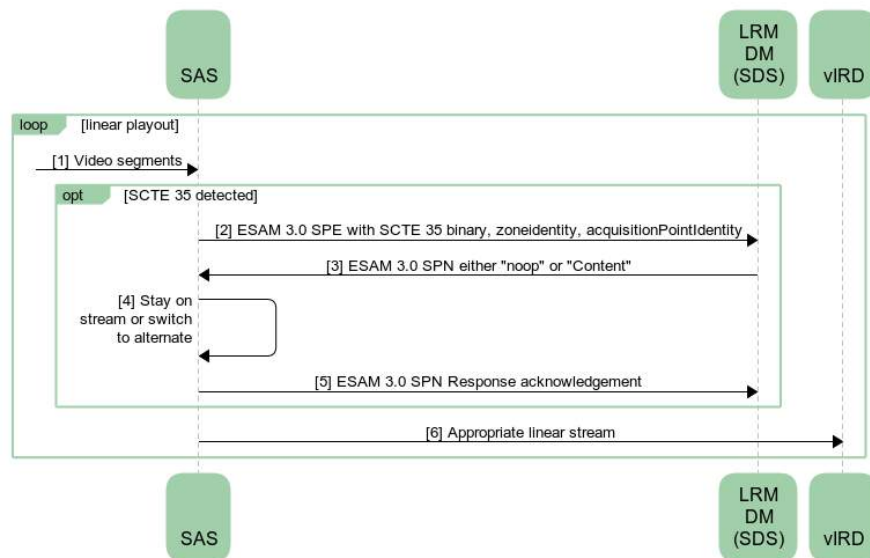
The ANSI/SCTE 35 standard describes an XML serialization of ANSI/SCTE 35 as shown in Figure 8 (with namespace declarations omitted for readability).

```
<SpliceInfoSection protocolVersion="0" ptsAdjustment="0" tier="4095"
xmlns="http://www.scte.org/schemas/35/2013a"
xmlns:ns2="http://www.thistech.com/schemas/scte35/1">
 <EncryptedPacket encryptionAlgorithm="0" cwIndex="255"/>
 <TimeSignal>
 <SpliceTime ptsTime="4151498432"/>
 </TimeSignal>
 <SegmentationDescriptor segmentationEventId="1"
segmentationEventCancelIndicator="false" segmentationTypeId="17" segmentNum="0"
segmentsExpected="0">
 <DeliveryRestrictions webDeliveryAllowedFlag="false"
noRegionalBlackoutFlag="false" archiveAllowedFlag="false" deviceRestrictions="0"/>
 <SegmentationUpid
segmentationUpidType="1">30303039353532303839</SegmentationUpid>
 </SegmentationDescriptor>
</SpliceInfoSection>
```

**Figure 8 - Example ANSI/SCTE 35 XML Serialization**

## 2.3. ANSI/SCTE 250

The “Real-time Event Signaling and Management API” (ESAM; ANSI/SCTE-250) protocol is sufficient for communication between a signal acquisition system (SAS) and signal decision system (SDS). It describes the format and content of the messages required by the SAS to obtain instructions from the decisioning service for handling in-band ANSI/SCTE 35 signals (Figure 9).



**Figure 9 - Traditional ESAM Exchange**

### 2.3.1. Message Transport

Messages are formatted and delivered in accordance with the CableLabs ESAM I03 specification. (CableLabs ESAM) The signal processing services are exposed as an HTTP RESTful endpoint. HTTP POST is the required request method and XML is the required message format, so the Content-Type header is set to application/xml. The SAS in this model is expected to return acknowledgements and processing status in the response.

### 2.3.2. Message Namespace

Table 2 shows the prefixes and namespaces that are used within the XML in ANSI/SCTE 35.

**Table 2 - Example ANSI/SCTE 35 Namespaces**

|           |                                                |
|-----------|------------------------------------------------|
| adi3      | urn:cablelabs:md:xsd:core:3.0                  |
| common    | urn:cablelabs:iptvservices:esam:xsd:common:1   |
| content   | urn:cablelabs:md:xsd:content:3.0               |
| manifest  | urn:cablelabs:iptvservices:esam:xsd:manifest:1 |
| offer     | urn:cablelabs:md:xsd:offer:3.0                 |
| po        | urn:cablelabs:md:xsd:placementopportunity:3.0  |
| signal    | urn:cablelabs:iptvservices:esam:xsd:signal:1   |
| signaling | urn:cablelabs:md:xsd:signaling:3.0             |
| terms     | urn:cablelabs:md:xsd:terms:3.0                 |
| title     | urn:cablelabs:md:xsd:title:3.0                 |
| xsi       | urn:cablelabs:iptvservices:esam:xsd:signal:1   |

### 2.3.3. Message Time

The CableLabs ESAM I03 specification requires that all object time values follow the ISO 8601 time format, hh:mm:ss, with an optional decimal fraction on the seconds component, e.g. 14:15:03.475. The

specification also dictates that times are to be provided as zero UTC offset. All times in this document are zero UTC offset and may contain fractional seconds, as per the specification.

### 2.3.4. ESAM Interface

The ESAM interface supports signal confirmation. The ESAM interface carries ANSI/SCTE 35 signaling information from an SAS to the SDS for in-band processing. The originating system generates a processing signal event and submits it using the defined XML payload to the SDS. The event parameters provide as much information as possible about the signal point. Examples of in-band signal points are a splice out/exit point indicated by an ANSI/SCTE 35 splice\_info\_section()/splice\_insert() command or an SCTE 35 segmentation\_descriptor associated with a SCTE 35 time\_signal() or splice\_null() command. Using the provided information, the SDS derives information about the event start and end point(s) with the corrected timestamp and/or ANSI/SCTE 35 point data. Finally, the SDS generates a response notification detailing how the acquisition system might condition the video content and provides auxiliary data for downstream usage.

### 2.3.5. In-band ESAM Message Exchange

A request via the SAS utilizes the SignalProcessingEvent message as a wrapper to the AcquiredSignal element. The acquisitionPointIdentity attribute is used to correlate the stream to the proper source in the signal processor. The content of the AcquiredSignal element is a ANSI/SCTE 35 descriptor and is sent as a binary payload. The binary payload is Base64 encoded.

## 2.4. Use Cases

### 2.4.1. Decision Request

Upon detection of a ANSI/SCTE 35 marker, the SAS will assemble the ESAM Signal Processing Event (SPE) (Figure 10). For the SPE, the acquisitionPointIdentity must be set to the source of the video stream (i.e., KXYZ). The acquisitionSignalID is a uuid established by the SAS for correlating the SPE with other responses for tracking and logging. The acquisitionTime like acquisitionPointIdentity is required as well and represents the time the linear stream stitcher (LSS) created the SPE message. The last required element is zoneIdentity used to provide the name of the virtual integrated receiver decoder (vIRD) or station instance for matching against the ANSI/SCTE 224 audiences (i.e., WXYZ). Because ESAM does not provide a mechanism to differentiate between vIRDs and zipcodes, the SDS will use any five-digit numbers as a ZIP type audience. The UTCPoint represent the UTC of the signal in the video stream. Finally, the BinaryData element contains the binary data from the ANSI/SCTE 35 marker found in the video stream.

```
<esam:SignalProcessingEvent xmlns:esam="urn:cablelabs:iptvservices:esam:xsd:signal:1"
 xmlns:core="urn:cablelabs:md:xsd:core:3.0"
 xmlns:md="urn:cablelabs:md:xsd:signaling:3.0"
 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
 xsi:schemaLocation="urn:cablelabs:iptvservices:esam:xsd:signal:1 OC-SP-ESAM-API-I03-
Signal.xsd">
 <esam:AcquiredSignal
 acquisitionPointIdentity="CH1"
 acquisitionSignalID="3c1e5cbe-648a-4d69-b740-570a99cb66d7"
 acquisitionTime="2018-04-05T00:00:04Z"
 zoneIdentity="KXYZ">
 <md:UTCPoint utcPoint="2018-04-04T20:00:06.281Z"/>
 <md:BinaryData
signalType="SCTE35">/DBKAAAAAAAAAP/wBQb+u6m/DgA0AjJDVUVJAAAAAX//AARgYLMbHjAwMDM1TUEwMDAwMDAwMzE0MDdUMD
QwNTE4MDAzMBABAB4s/do=</md:BinaryData>
```



```
</esam:AcquiredSignal>
</esam:SignalProcessingEvent>
```

**Figure 10 - Example ESAM SignalProcessingEvent**

### 2.4.2. Decision Response

Once the SDS receives the ESAM decision request, it will execute the SCTE 224 logic described above and return to the SAS the action required based on the ViewingPolicies. The response will contain the action of the source to switch to or “noop,” indicating the SAS does not need to switch from the source it is currently on.

Figure 11 shows the response the SAS will receive for no switch:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<signal:SignalProcessingNotification
 xmlns:adi3="urn:cablelabs:md:xsd:core:3.0"
 xmlns:signaling="urn:cablelabs:md:xsd:signaling:3.0"
 xmlns:signal="urn:cablelabs:iptvservices:esam:xsd:signal:1"
 xmlns:manifest="urn:cablelabs:iptvservices:esam:xsd:manifest:1"
 xmlns:ns5="http://www.cablelabs.com/namespaces/metadata/xsd/confirmation/2"
 xmlns:common="urn:cablelabs:iptvservices:esam:xsd:common:1"
 xmlns:content="urn:cablelabs:md:xsd:content:3.0"
 xmlns:offer="urn:cablelabs:md:xsd:offer:3.0"
 xmlns:po="urn:cablelabs:md:xsd:placementopportunity:3.0"
 xmlns:terms="urn:cablelabs:md:xsd:terms:3.0"
 xmlns:title="urn:cablelabs:md:xsd:title:3.0">
 <common:StatusCode classCode="0"/>
 <signal:ResponseSignal action="noop" acquisitionPointIdentity="CH1" acquisitionSignalID="27c3bf78-
3f65-4f03-ac44-4f61de204991" acquisitionTime="2018-05-18T15:53:03Z">
 <signaling:UTCPPoint utcPoint="2018-05-18T15:53:07.184Z"/>
 <signaling:BinaryData
signalType="SCTE35"/>DBKAAAAAAAAAAP/wBQb+m3lcMAA0AjJDVUVJAAAAAX//AAG2TxcBHjAwMDM0TUEwMDAwMDAwMzI2ODJUMD
UxODE4MTUwMAEBADzcPCU=</signaling:BinaryData>
 </signal:ResponseSignal>
</signal:SignalProcessingNotification>
```

**Figure 11 - Example ESAM SignalProcessingNotification Noop**

Figure 12 shows the response the SAS will receive for a switch:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<signal:SignalProcessingNotification
 xmlns:adi3="urn:cablelabs:md:xsd:core:3.0"
 xmlns:signaling="urn:cablelabs:md:xsd:signaling:3.0"
 xmlns:signal="urn:cablelabs:iptvservices:esam:xsd:signal:1"
 xmlns:manifest="urn:cablelabs:iptvservices:esam:xsd:manifest:1"
 xmlns:ns5="http://www.cablelabs.com/namespaces/metadata/xsd/confirmation/2"
 xmlns:common="urn:cablelabs:iptvservices:esam:xsd:common:1"
 xmlns:content="urn:cablelabs:md:xsd:content:3.0"
 xmlns:offer="urn:cablelabs:md:xsd:offer:3.0"
 xmlns:po="urn:cablelabs:md:xsd:placementopportunity:3.0"
 xmlns:terms="urn:cablelabs:md:xsd:terms:3.0"
 xmlns:title="urn:cablelabs:md:xsd:title:3.0">
 <common:StatusCode classCode="0"/>
 <signal:ResponseSignal action="noop" acquisitionPointIdentity="CH1" acquisitionSignalID="27c3bf78-
3f65-4f03-ac44-4f61de204991" acquisitionTime="2018-05-18T15:53:03Z">
 <signaling:UTCPPoint utcPoint="2018-05-18T15:53:07.184Z"/>
 <signaling:BinaryData
signalType="SCTE35"/>DBKAAAAAAAAAAP/wBQb+m3lcMAA0AjJDVUVJAAAAAX//AAG2TxcBHjAwMDM0TUEwMDAwMDAwMzI2ODJUMD
UxODE4MTUwMAEBADzcPCU=</signaling:BinaryData>
```

```

 <signal:AlternateContent altContent="true" altContentIdentity="CH5" zoneIdentity="KXYZ"/>
 </signal:ResponseSignal>
</signal:SignalProcessingNotification>

```

**Figure 12 - Example ESAM SignalProcessingNotification for a Switch**

### 2.4.3. SAS Response Acknowledgement

Upon receiving the SDS SPN response, the SAS will return an ESAM acknowledgement indicating the action taken by the SAS in the notes section. For the acknowledgement there are no specific requirements for the StatusCode::classCode for the ESAM-endpoint. The status code is only logged, but not used directly. For the Note element, there are also no specific requirements for its values. SDS logs this value, and currently SDS sees three common responses: “Successfully switched to CHXX,” “Cannot switch, because already scheduled to switch to that channel,” and “No switch, staying on CHXX” (Figure 13).

```

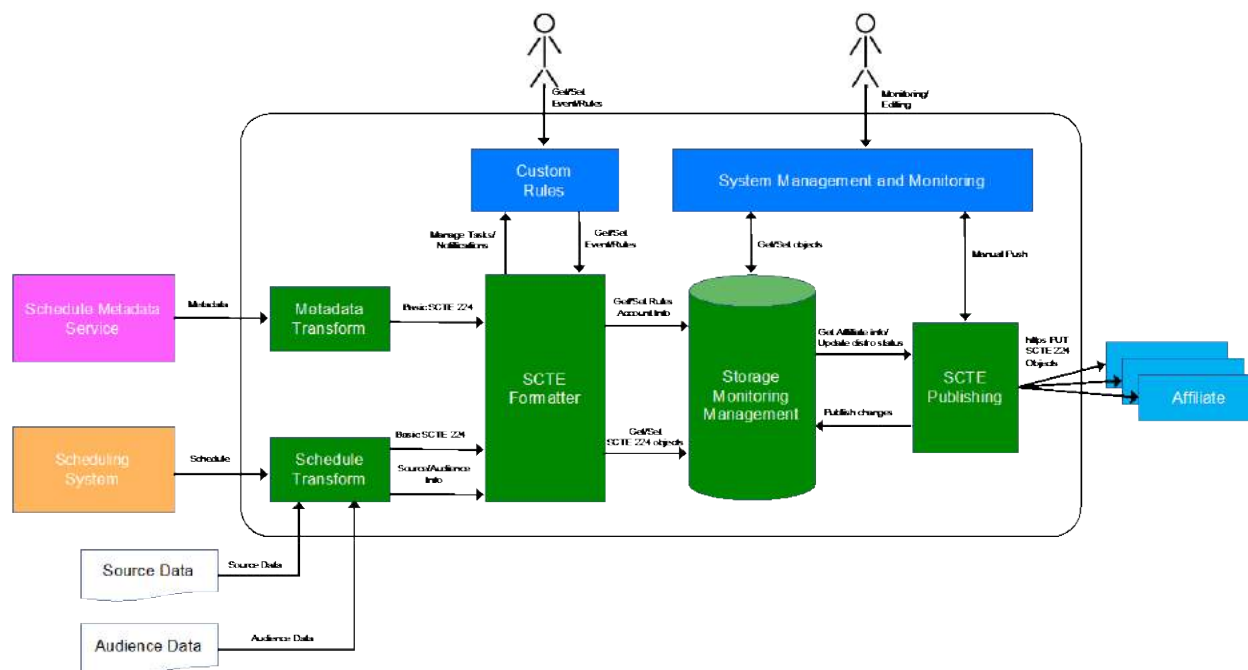
<?xml version="1.0" encoding="UTF-8"?>
<esam:ProcessStatusNotification acquisitionPointIdentity="KXYZ_PHILADELPHIA" acquisitionSignalID=
"3c1e5cbe-648a-4d69-b740-570a99cb66d7" xmlns:core="urn:cablelabs:md:xsd:core:3.0"
xmlns:esam="urn:cablelabs:iptvservices:esam:xsd:common:1" xmlns:md="urn:cablelabs:md:xsd:content:3.0"
xmlns:sig="urn:cablelabs:md:xsd:signaling:3.0" xmlns:xml="http://www.w3.org/XML/1998/namespace"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="urn:cablelabs:iptvservices:esam:xsd:common:1 OC-SP-ESAM-API-I03-Common.xsd ">
 <esam:StatusCode classCode="0" detailCode="0">
 <core:Note>No switched needed source is already active : CH5 </core:Note>
 </esam:StatusCode>
</esam:ProcessStatusNotification>

```

**Figure 13 - Example ESAM SignalProcessingNotification Acknowledgement**

## 3. Distribution Workflow

With the ongoing consumer demands for higher quality video, as well as the new C-band rules from the FCC, content providers need to strongly consider a transition to non-satellite delivery. The workflow is similar whether there is a broadcaster distributing their 12-16 hours of daily content to an affiliate station, or if it is an affiliate station distributing their 24 hours of content to a distribution partner. In both cases there is content that needs to get to a distribution partner to be played out to a consumer. This distribution workflow can be done with or without in-band signaling. This section will discuss the differences presented by the presence or absence of in-band signaling and how either situation can be handled.



**Figure 14 – Detailed Workflow**

### 3.1. Transition with Automation

Many content providers use in-band signals even in the satellite feeds they are distributing today. This is most often the case for the full channel content being sent to multiprogram video program distributors (MVPD) or an OTT streaming service. This usually means that the content provider has an automation system capable of creating ANSI/SCTE 104 messages that are used to transfer signal insertion instructions from an automation system and convey those markings to the compression system. This ultimately results in creating the corresponding ANSI/SCTE 35 in-band markers for distribution in the video itself. If that is the case, and all program starts and ends are identified with ANSI/SCTE 35 binary signals, then in-band signals need not be injected into the video. Note that section 3.2 “Transition without Automation” describes the case where signals are missing and need to be injected.

#### 3.1.1. Signaling

When in-band signals are sufficient to declare the content start and end, the burden on the content provider is to extract or correlate those ANSI/SCTE 35 UPIDs with the ANSI/SCTE 224 ESNI that is generated. These UPIDs must be extracted from the ANSI/SCTE 35 binary and placed in the “MatchSignal” element of the ANSI/SCTE 224 ESNI MediaPoint that goes along with that event. So, for example, if a system uses the show’s Tribune Media Service identifier (TMS ID) as the Upid to uniquely identify a program, then that TMS ID can be used by the ANSI/SCTE 224 generate platform to populate the XPath matching in the MediaPoint. Figure 15 shows an example of the MatchSignal portion of a ANSI/SCTE 224 MediaPoint.

```

<MatchSignal xmlns="http://www.scte.org/schemas/224" match="ANY">
 <Assert>
 /SpliceInfoSection/SegmentationDescriptor[@segmentationTypeId=16]/SegmentationUpid[@segmentationUpidType=1 and contains(text(),'EP001786121816')]
 </Assert>
</MatchSignal>

```

### Figure 15 - Example MediaPoint MatchSignal

Although automation inserts ANSI/SCTE 35 signals in the video stream, downstream automation needs to get Upid information (such as “EP001786121816” from the example above) to the ANSI/SCTE 224 ESNI creation system.

#### 3.1.2. Additional Instruction

At this point the in-band signaling information is in a MediaPoint’s MatchSignal element. On playout, if the SAS received the in-band signal, the SDS would match on a MediaPoint. The SDS replies with the “Action” to be taken by the acquisition system, based on the audience of interest. This is done in ANSI/SCTE 224 by using “Apply” policy to activate a policy. The Policy then points to the ViewingPolicy, which has the real functionality of bringing an audience together with an action.

For a broadcast solution, where a content provider is selecting a feed for a particular show, the provider can identify the desired audience, for example by ZIP code or audience type. If the SAS identity matches the audience then the action applies. The action could indicate the desired content, such as from an alternate channel source, or a URL from which the content could be retrieved. With this model, a content provider could tell one partner to play content from one channel and another partner to play content from another.

For example, a content provider of a broadcast network can tell its Eastern time zone partners to pull content from one feed, while telling its Central time zone partners to pull from a different feed. Alternatively, for a live event, the content provider could tell all partners to source from the same live feed.

Likewise, a content provider distributing a channel could tell one partner to black out or display a slate during a particular program, while having another region or provider play the content on the current feed. This situation is useful for sports blackouts or when content rights vary between regions.

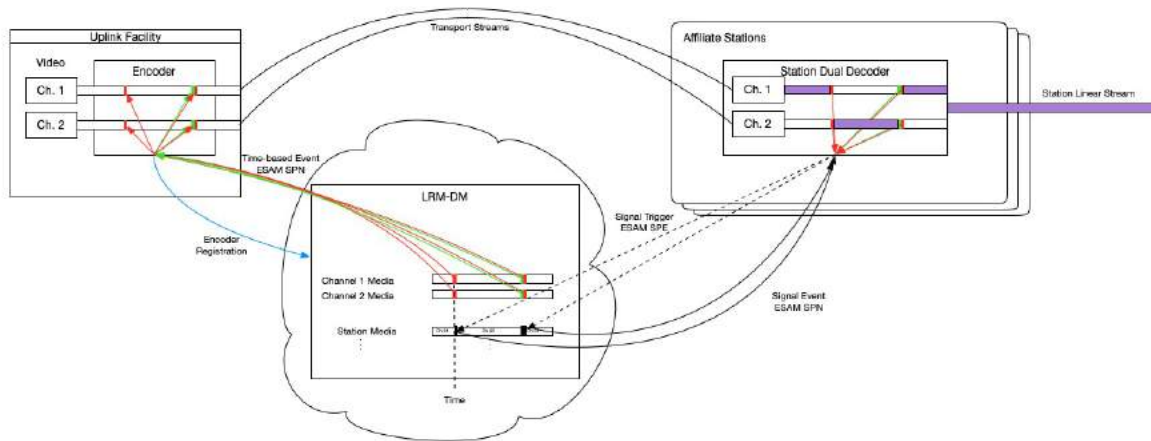
### 3.2. Transition without Automation

In many cases, content providers do not have sufficient in-band signaling in their video feeds, possibly because satellite feeds rarely needed detailed signaling, or because most or all of the controls are done by integrated receiver/decoders (IRDs). Another situation is that OTT feeds, or higher quality feeds that are sourced terrestrially, are not subject to the same workflow. Additionally, the automation system driving the satellite delivery might lack signaling features whatsoever. SCTE 224 ESNI accommodates for these cases and still allows signals to be inserted without the corresponding automation.

#### 3.2.1. Signal Insertion

The ANSI/SCTE 224 ESNI standard provides actions in the ViewingPolicy to enable signal insertion. In this sense, the content provider delivers schedule information to the generator; the generator instructs the SAS to inject signals when needed. In this case, generator creates two sets of instructions in the system. The first set of Media contains MediaPoints for each show start and stop, but the action is a “SignalPointInsertion” action. The second set of Media are the same MediaPoints for the shows, but this Media contains the MatchSignal logic to connect to the signals that are being inserted. This second set of Media is for distribution to execute specific actions.

Figure 16 depicts the first Media, marking the distributed video for edge execution, and the second Media being used at the edge to execute the content provider's rights for the content.



**Figure 16 - Workflow for Signal Insertion**

When the scheduling system sends the ANSI/SCTE 224 ESNI generator the schedule information, the generator has to know to generate two different Media elements. The start and end times are the same for both Media elements. The difference is the Policy that the MediaPoints point to, and that the MatchSignal element is only in one of the Media elements.

### **3.2.1.1. Signal Insertion Media**

The signal insertion Media contains the audience of the channel into which the signals will be inserted. If the event will not be triggered by in-band signals, the insertion will be initiated by the SDS itself. These SDS-initiated ESAM communications are referred to as “unsolicited messages.” In these cases, the MediaPoints in the containing Media element will have a matchtime attribute that represents the precise time that the event will occur, and the signal should be inserted.

In many SDS systems, a time-based, unsolicited message is coupled with an ANSI/SCTE 250 ESAM registration by the SAS. This registration tells the SDS which endpoint to send the unsolicited messages. It also usually contains lead-time information, about when to send the instructions, and, finally, the desired audience. For example, if a show starts on feed X at 2:00 PM, and an SAS registration is interested in receiving the signaling for feed X with a one minute lead time, then the SDS sends an unsolicited ESAM SignalProcessingNotification message at 1:59 PM to the SDS endpoint. This gives the SAS the time necessary to insert the signal at the proper time in the video.

Once this in-band signal has been inserted, the SAS will match it on the execution side using the other Execution Media.

### **3.2.1.2. Distribution Media**

The second Media generated from the show's schedule has the same start and ends times, but in this Media the ESNI generator creates the MatchSignal elements and inserts the signal information used

above. It also can populate the show's asset distribution interface (ADI) information if the content provider so desires, since this Media is used on the edge by the distribution partner.

### **3.2.1. Additional Instruction**

This Distribution Media then follows the same execution path that the “Automation” section described, where the SAS detects the in-band signal and sends it to the SDS for a decision about what to do at the point in the video feed. The SDS then responds with the appropriate action for the audience represented by the SAS.

## **4. Solution**

At Comcast Technology Solutions, we created a solution that bridges the creation, management and distribution process of ANSI/SCTE 224 events for both broadcast and OTT environments. The service automates the manual processes and takes full advantage of the ANSI/SCTE 224 specification – not just by standardizing the formatting of metadata, but also by providing a software-based service that gives both programmers and operators unprecedented control of – and visibility into – their respective workflows. As more variables are introduced into the content delivery mix, whether they be new content types or new viewing policy actions, we will always need a flexible strategy that can deliver extraordinary experiences across every screen. Providers can simplify processes and utilize better content management tools, and then apply them to what matters most: creating an experience that turns viewers into ardent and vocal fans.

## **5. Additional thoughts**

### **5.1. OTT Complexity**

Complex viewing rights have become an increasing challenge for content delivery. No matter what the specific rights are for a program -- whether they pertain to regional networks, league- or match-specific arrangements, geography or user permissions -- it has become crucial that providers are precisely conveying the policies necessary to allow for authorized playback. Digital distribution partners require these rights in order to distribute the client's content outside the home viewing territory. ANSI/SCTE 224 stands out as the best way to deliver this information. The new metadata management approach is now an integral part of operational workflows and is managed by an internal operations team.

The solution supports both MatchTime-based event rules and MatchSignal-based event rules. This means a full functional adaptation of ANSI/SCTE 35 segmentation UPID, or signal ID's, at the ProgramStart and ProgramEnd boundaries of a MediaPoint, and is able to map the ANSI/SCTE 35 ID to the MatchSignal ID in the ANSI/SCTE 224 event. A clean transition is ensured when alternate content rules are applied to the linear feed. When an event window runs beyond the scheduled airing time, the alternate content execution is based on the segmentation UPID identified in the video feed, matching the UPID that is in the ANSI/SCTE 224 event. This allows for precision control from the content provider, while maintaining a positive user experience for the viewer.

Linear rights data can now be distributed as many as 14 days ahead of the airing of the event, and real-time distribution solves for event changes, such as a baseball game that goes into extra innings. As a broadcaster that distributes hundreds of regional sports networks, managed at a national level while being distributed by digital distribution partners or OTT services, solving for unscheduled event changes is important. Using a push model, linear programming events are distributed to the metadata management tool. Operations teams now have visibility into the program schedule for all channels, because that data is

normalized into ANSI/SCTE 224 and then distributed to all digital distribution partners. Through a desktop console, operations teams can:

- Make real-time changes to linear events
- View programming schedules
- Edit and create audience definitions
- Perform policy updates
- Validate decision logic and audit activity

Following the ANSI/SCTE 224 standard, all events are distributed using XML over HTTP to designated URL endpoints. Messages have an option to be signed using an assigned key value, and have a configurable frequency as well as a retry rate. The messages can also be individually targeted to each distribution partner and filtered, so each partner gets only what they need.

As an organization among the early adopters of ANSI/SCTE 224, it's worth noting that several of the largest broadcasters, at the forefront of the industry, are using these extant standards and resultant tools to take the lead on how linear rights metadata can and should be delivered.

## **5.2. EPG Usage**

An additional benefit of transitioning to ANSI/SCTE 224 ESNI is ease in sharing data used to populate EPGs. Once broadcasters have an OTT schedule completely decorated in ESNI, then that schedule can be shared with an OTT partner, who can process the information in a separate workflow for the EPG. For example, if the ESNI schedule is populated for seven days in the future, then an OTT distribution partner can also populate its EPG into the future. Unlike the generic program guide information, an ESNI schedule contains alternate content information. Thus, the partner can present the subscriber with a highly accurate, tailored EPG. They may then look ahead in the schedule and see the “real” guide for the specific experience.

Likewise, if the content provider is making changes to the playout schedule, the corresponding ANSI/SCTE 224 ESNI schedule changes immediately. It flows to the distribution partner instantaneously, and faster than third-party services that manage program guide data.

## **6. Conclusion**

As broadcasters adapt their access to and use of the C-band spectrum for traditional video distribution, concurrent with the FCC's reallocation of the lower 280 MHz of that band for 5G wireless service, many are opting to shift to terrestrial, fiber-based delivery. It's a compelling shift, with benefits that include lower costs, lower latencies, fewer bandwidth constraints, to name but a few.

Any transition to terrestrial video delivery over fiber must occur with minimal impact to current customer workflows. The risks of transition certainly include disruption of primary service, but also of signaling for advertising, local contributions, and alternate programming, including blackouts. The benefits include additional OTT distribution. Certain hybrid use cases increase disruption risk. For example, national feeds arriving by satellite can signal a local contribution arriving over fiber, or vice versa.

The SCTE DDVS maintains a family of standards that, when implemented together, help support a seamless transition from satellite to fiber. These standards can be applied, in real-world situations, to enable a transition to terrestrial, fiber-based delivery with minimal impacts to workflows.

This paper detailed an architecture and specifications for a linear content signaling and decision system that is useful for implementing the intricacies of rights management, affiliate programming, and satellite delivery replacement. It detailed the use of ANSI/SCTE-35, ANSI/SCTE-224 (ESNI) and ANSI/SCTE-250 (ESAM) into an architecture useable by on-premise and cloud-based systems.

Comcast Technology Solutions has developed and deployed this system for hundreds of channels and thousands of decisions per week. Implementing this linear system enables the seamless transition of fully featured linear channels, from satellite to the Internet.



## Abbreviations

|       |                                                     |
|-------|-----------------------------------------------------|
| ANSI  | American National Standards Institute               |
| API   | application programmer interface                    |
| ADI   | asset distribution interface                        |
| dMVPD | digital multichannel video programming distributors |
| DVS   | [SCTE] Digital Video Subcommittee                   |
| EPG   | electronic program guide                            |
| ESAM  | Event Signaling and Management                      |
| ESNI  | Event Scheduling and Notification Interface         |
| FCC   | Federal Communications Commission                   |
| HTTP  | Hypertext Transfer Protocol                         |
| IRD   | integrated receiver decoder                         |
| ISBE  | International Society of Broadband Experts          |
| ISO   | International Standards Organization                |
| LSS   | linear stream stitcher                              |
| MPEG  | Moving Pictures Experts Group                       |
| MVPD  | multichannel video programming distributors         |
| OTT   | over-the-top                                        |
| REST  | Representational State Transfer                     |
| SAS   | signal acquisition service                          |
| SCTE  | Society of Cable Telecommunications Engineers       |
| SDS   | signal decisioning service                          |
| SPE   | signal processing event                             |
| SPN   | signal processing notification                      |
| TMS   | Tribune Media Service                               |
| UPID  | unique program identifiers                          |
| URL   | Uniform Resource Locator                            |
| UTC   | Universal Time Coordinated                          |
| vIRD  | virtual integrated receiver decoder                 |
| XML   | Extensible Markup Language                          |

## Bibliography & References

CableLabs ESAM, Real-time Event Signaling and Management API, 2013. [https://scte-cms-resource-storage.s3.amazonaws.com/SCTE\\_35\\_OC-SP-ESAM-API-I03-131025.pdf](https://scte-cms-resource-storage.s3.amazonaws.com/SCTE_35_OC-SP-ESAM-API-I03-131025.pdf).

FCC Report and Order on C-band. Expanding Flexible Use of the 3.7 to 4.2 GHz Band  
<https://docs.fcc.gov/public/attachments/FCC-20-22A1.pdf>

Reardon, M. February 6, 2020. FCC reaches deal with satellite industry to free up more 5G spectrum, CNet, retrieved from: <https://www.cnet.com/news/fcc-reaches-deal-with-satellite-industry-to-free-up-more-5g-spectrum/>

SCTE 224 ESNI, Event Scheduling and Notification Interface 2018r1. [https://scte-cms-resource-storage.s3.amazonaws.com/Standards/ANSI\\_SCTE%20224%202018r1.pdf](https://scte-cms-resource-storage.s3.amazonaws.com/Standards/ANSI_SCTE%20224%202018r1.pdf)

SCTE 35, Digital Program Insertion Curing Message for Cable, 2019. [https://scte-cms-resource-storage.s3.amazonaws.com/ANSI\\_SCTE-35-2019a-1582645390859.pdf](https://scte-cms-resource-storage.s3.amazonaws.com/ANSI_SCTE-35-2019a-1582645390859.pdf)

# Appendix A

## Example ANSI/SCTE 224 Media Object

```
<Media xmlns="http://www.scte.org/schemas/224"
 id="xyz/media/WXYZ"
 description="WXYZ"
 lastUpdated="2020-03-24T17:49:45.000Z">
 <MediaPoint xmlns="http://www.scte.org/schemas/224"
 id="xyz/media/WXYZ/resident"
 description="WXYZ"
 lastUpdated="2020-03-24T17:49:45.000Z"
 effective="2020-03-23T18:00:00.000Z"
 expires="2020-04-01T08:30:00.000Z">
 <Apply xmlns="http://www.scte.org/schemas/224">
 <Policy xmlns="http://www.scte.org/schemas/224"
 xmlns:xlink="http://www.w3.org/1999/xlink"
 xlink:href="xyz/policy/any/slate"></Policy>
 </Apply>
 </MediaPoint>
 <MediaPoint xmlns="http://www.scte.org/schemas/224"
 id="xyz/media/WXYZ/program/3E264A41-8BE7-41D9-8FAF-2072EFA0A866/start"
 description="Show 1"
 lastUpdated="2020-03-24T17:49:45.000Z"
 effective="2020-03-23T17:30:00.000Z"
 expires="2020-03-23T19:30:00.000Z"
 source="3230">
 <AltID xmlns="http://www.scte.org/schemas/224">EP002830833983</AltID>
 <Metadata xmlns="http://www.scte.org/schemas/224">
 <MetadataDetail xmlns="http://ctsrmm.com/ctsesni"
 name="ScheduledAiringId"
 type="string"
 provider="XYZ">3E264A41-8BE7-41D9-8FAF-2072EFA0A866</MetadataDetail>
 <MetadataDetail xmlns="http://ctsrmm.com/ctsesni"
 name="ScheduledStart"
 type="string"
 provider="XYZ">2020-03-23T18:00:00.000Z</MetadataDetail>
 <MetadataDetail xmlns="http://ctsrmm.com/ctsesni"
 name="ScheduledEnd"
 type="string"
 provider="XYZ">2020-03-23T19:00:00.000Z</MetadataDetail>
 <MetadataDetail xmlns="http://ctsrmm.com/ctsesni"
 name="StartOver" type="string"
 provider="XYZ">false</MetadataDetail>
 <MetadataDetail xmlns="http://ctsrmm.com/ctsesni"
 name="LookBack"
 type="string"
 provider="XYZ">false</MetadataDetail>
 </MetadataDetail>
 <ADI3 xmlns="http://www.scte.org/schemas/236/2017/core">
 <Asset xmlns:XMLSchema-instance="http://www.w3.org/2001/XMLSchema-instance"
 XMLSchema-instance:type="title:TitleType"
 uriId="xyz.com/Title/3E264A41-8BE7-41D9-8FAF-2072EFA0A866"
 startDateTime="2020-03-27T14:14:40Z"
 endDateTime="2020-03-27T14:30:16Z"
 lastModifiedDateTime="2020-03-27T15:04:46.744Z">
 <Provider>XYZ</Provider>
 </ADI3>
 <Ext>
 <App_Data Name="Season_Number" Value="1"></App_Data>
 <App_Data Name="Episode_Number" Value="S1:E14"></App_Data>
 <App_Data Name="Series_Name" Value="Show 1"></App_Data>
 <App_Data Name="Episode_Title" Value="Show 1"></App_Data>
 <App_Data Name="Season_Title" Value="Show 1"></App_Data>
 <App_Data Name="Category" Value="Drama"></App_Data>
 <App_Data Name="Is_Live" Value="N"></App_Data>
 <App_Data Name="DVS" Value="N"></App_Data>
 </Ext>
 </MediaPoint>
</Media>
```

```

 <LocalizableTitle xmlns="http://www.scte.org/schemas/236/2017/title">
 <TitleBrief xmlns="http://www.scte.org/schemas/236/2017/title">Maury</TitleBrief>
 <TitleMedium xmlns="http://www.scte.org/schemas/236/2017/title">Maury</TitleMedium>
 <TitleLong xmlns="http://www.scte.org/schemas/236/2017/title">Maury</TitleLong>
 <SummaryShort xmlns="http://www.scte.org/schemas/236/2017/title">Funny
Story</SummaryShort>
 </LocalizableTitle>
 <Rating xmlns="http://www.scte.org/schemas/236/2017/title"
 ratingSystem="TV">TV-Y7-FV</Rating>
 <IsClosedCaptioning xmlns="http://www.scte.org/schemas/236/2017/title">
 true</IsClosedCaptioning>
 <DisplayRunTime xmlns="http://www.scte.org/schemas/236/2017/title">
 00:15</DisplayRunTime>
 <Year xmlns="http://www.scte.org/schemas/236/2017/title">2019</Year>
 <ShowType xmlns="http://www.scte.org/schemas/236/2017/title">Series</ShowType>
 </Asset>
 <Asset xmlns:XMLSchema-instance="http://www.w3.org/2001/XMLSchema-instance"
 XMLSchema-instance:type="content:MovieType"
 uriId="xyz.com/Asset/3E264A41-8BE7-41D9-8FAF-2072EFA0A866"
 providerVersionNum="10"
 internalVersionNum="0"
 creationDateTime="2020-03-26T22:20:19.762Z"
 startDateTime="2020-03-27T14:14:40Z"
 endDateTime="2020-03-27T14:30:16Z"
 lastModifiedDateTime="2020-03-27T15:04:46.744Z">
 <Provider>XYZ</Provider>
 <AudioType xmlns="http://www.scte.org/schemas/236/2017/content">Dolby 5.1</AudioType>
 <Language xmlns="http://www.scte.org/schemas/236/2017/content"
 bitStreamMode="0">en</Language>
 </Asset>
</ADI3>
</Metadata>
<Apply xmlns="http://www.scte.org/schemas/224">
 <Policy xmlns="http://www.scte.org/schemas/224"
 xmlns:xlink="http://www.w3.org/1999/xlink"
 xlink:href="xyz/policy/WXYZ.all/3230">
 </Policy>
</Apply>
<MatchSignal xmlns="http://www.scte.org/schemas/224" match="ANY">
 <Assert xmlns="http://www.scte.org/schemas/224">
 /SpliceInfoSection/SegmentationDescriptor[@segmentationTypeId=16]/
 SegmentationUpid[@segmentationUpidType=1 and contains(text(),'EP002830833983')]/</Assert>
 <Assert xmlns="http://www.scte.org/schemas/224">
 /SpliceInfoSection/SegmentationDescriptor[@segmentationTypeId=1]/
 SegmentationUpid[@segmentationUpidType=1 and contains(text(),'EP002830833983')]/</Assert>
 </MatchSignal>
</MediaPoint>
<MediaPoint xmlns="http://www.scte.org/schemas/224"
 id="xyz/media/WXYZ/program/3E264A41-8BE7-41D9-8FAF-2072EFA0A866/end"
 description="Show 1"
 lastUpdated="2020-03-24T17:49:45.000Z"
 effective="2020-03-23T18:30:00.000Z"
 expires="2020-03-23T19:30:00.000Z"
 source="3230">
 <Remove xmlns="http://www.scte.org/schemas/224">
 <Policy xmlns="http://www.scte.org/schemas/224"
 xmlns:xlink="http://www.w3.org/1999/xlink"
 xlink:href="xyz/policy/WXYZ.all/3230">
 </Policy>
 </Remove>
 <MatchSignal xmlns="http://www.scte.org/schemas/224" match="ANY">
 <Assert xmlns="http://www.scte.org/schemas/224">
 /SpliceInfoSection/SegmentationDescriptor[@segmentationTypeId=17]/
 SegmentationUpid[@segmentationUpidType=1 and contains(text(),'EP002830833983')]/</Assert>
 <Assert xmlns="http://www.scte.org/schemas/224">
 /SpliceInfoSection/SegmentationDescriptor[@segmentationTypeId=1]/
 SegmentationUpid[@segmentationUpidType=1 and
 not(contains(text(),'EP002830833983'))]/</Assert>
 </MatchSignal>
 </MediaPoint>

```

</Media>

# **Gridmetrics™ Data Provide Insights and Improve Situational Awareness of the Electric Power Grid**

## **Powering 10G: What It Takes & How to Do It**

A Technical Paper prepared for SCTE•ISBE by

**Robert Cruickshank, Ph.D.**

Consultant

Cable Television Laboratories / National Renewable Energy Laboratory

132 Cruickshank Rd # 269, Big Indian, NY 12410

+1-703-568-8379

rfciii@cruickshank.org

**Nicolas Metts**

Lead Software Engineer / Data Scientist

Cable Television Laboratories

858 Coal Creek Cir, Louisville, CO 80027

+1-303-661-9100

n.metts@cablelabs.com

**Paul Schauer**

OSS Architect

Comcast Cable Communications

183 Inverness Drive W, Englewood, CO 80111

+1-303-372-1215

paul\_schauer@comcast.com

**Curtis Snyder**

Consultant

11 Coopertown Road, Haverford, PA 19041

+1-215-341-2821

laurence.snyder@gmail.com

# Table of Contents

| <b>Title</b>                                             | <b>Page Number</b> |
|----------------------------------------------------------|--------------------|
| 1. Introduction .....                                    | 3                  |
| 2. Fundamental Cable Architecture .....                  | 4                  |
| 3. Normal Grid Operating Voltages.....                   | 5                  |
| 4. High Voltages .....                                   | 7                  |
| 5. Low Voltages .....                                    | 7                  |
| 6. Fluctuating Voltages.....                             | 8                  |
| 7. Outages.....                                          | 9                  |
| 8. Cable TV Broadband Sensor and Network Advantages..... | 13                 |
| 9. Conclusion .....                                      | 14                 |
| Abbreviations.....                                       | 15                 |
| Bibliography & References .....                          | 15                 |

## List of Figures

| <b>Title</b>                                                                        | <b>Page Number</b> |
|-------------------------------------------------------------------------------------|--------------------|
| Figure 1 – Cable tree-and-branch architecture .....                                 | 4                  |
| Figure 2 - Observed voltages in Gridmetrics national data set.....                  | 5                  |
| Figure 3 - Observed voltages in Gridmetrics national data set.....                  | 6                  |
| Figure 4 - Example of high voltage in the Gridmetrics national data set. ....       | 7                  |
| Figure 5 - Example of low voltage in the Gridmetrics national data set. ....        | 8                  |
| Figure 6 - Example of fluctuating voltage in the Gridmetrics national data set..... | 9                  |
| Figure 7 - Status of the secondary distribution grid prior to the outage.....       | 11                 |
| Figure 8 - Status of the secondary distribution grid during the outage.....         | 12                 |
| Figure 9 - Example of service near the end of the outage.....                       | 13                 |

## List of Tables

| <b>Title</b>                                              | <b>Page Number</b> |
|-----------------------------------------------------------|--------------------|
| Table 1 - Normal operating voltage limits [Volts AC]..... | 5                  |

# 1. Introduction

The newly available Gridmetrics data platform dramatically improves the situational awareness of the electric power grid worldwide by providing an entirely new and independent view of real-time, localized power conditions. Gridmetrics data deliver insights enabled by the existing cable broadband infrastructure that measures, monitors, and tracks the availability and stability of voltage and other power measures in the last mile of the grid known as the secondary distribution network. Voltage measurements can reflect stresses on the grid and provide an early warning of unstable, unsafe, or inefficient operating conditions. Insights include specific locations on the grid where voltage observations are abnormally high, low or fluctuating, and provide much-needed visibility to the vast majority of the grid, which heretofore has been sensor starved. Analysis reveals voltage anomalies and outages that likely impact customer experience, safety, and equipment longevity.

Situational awareness of the electric power grid is gaining in importance with the increasing number of power generators, powered devices, and power infrastructure failures. In the United States, 200,000 miles of well-instrumented high-voltage transmission lines make up the grid core, or backbone. However, the lower voltage, less-instrumented local distribution lines that connect the remaining 96.5% of the grid to end users, account for 5.5 million miles, with limited network visibility in the last mile.<sup>1</sup>

The rapidly growing penetration of electric vehicles (EVs), distributed energy resources (DERs), such as solar generation and battery storage, and connected buildings are creating new end-use dynamics and 2-way power flows that the existing grid was not designed to accommodate. These changes at the “grid edge” in the distribution network necessitate improving situational awareness to manage operations and detect grid stresses. State and local mandates to increase penetration of EVs and DERs, while critical to address the global climate crisis, further accelerate the changing grid edge dynamics.

Despite a mean life expectancy of 65 years, the average grid infrastructure element is 68 years old, with some elements well past the century mark. Unfortunately, we have already witnessed the effects of this aging infrastructure, through decreased reliability and deadly wildfires due to mechanical failures of the hooks that hold charged wires.<sup>2,3</sup> It is estimated that the U.S. power grid requires \$51 billion of annual maintenance and capital upgrades (up 57% from 2016 to 2018) to improve reliability and catch-up on deferred maintenance.<sup>4</sup> Using information on the locations of voltage anomalies, provided through the Gridmetrics platform, can help electric utility managers to more quickly and efficiently target maintenance and upgrades on problem areas, increasing grid reliability and, ultimately, saving lives.

Gridmetrics is an early-stage project being incubated at CableLabs that connects and overlays a vast high-speed private communications network to supply as-delivered performance, insights, and knowledge about the power distribution grid. A new, widely available, and rapidly growing

---

<sup>1</sup> <https://www.scientificamerican.com/article/what-is-the-smart-grid/>

<sup>2</sup> <https://www.wsj.com/articles/pg-e-knew-for-years-its-lines-could-spark-wildfires-and-didnt-fix-them-11562768885>

<sup>3</sup> <https://www.wsj.com/articles/this-old-metal-hook-could-determine-whether-pg-e-committed-a-crime-11583623059>

<sup>4</sup> <https://www.utilitydive.com/news/aging-grids-drive-51b-in-annual-utility-distribution-spending/528531/>




dataset of voltage observations from hundreds of thousands of in-service cable broadband power sensors provides 5-minute updates, insights, and situational awareness of the distribution grid. Work from early Gridmetrics participants has shown a path to significantly more frequent updates, providing an increasingly real-time big data stream that represents multiple orders of magnitude improvement on existing distribution grid monitoring solutions. Power industry experts have estimated that utilities would have to invest billions of dollars over multiple decades to replicate the existing Gridmetrics sensor platform, which at the same time would have already evolved to superior higher fidelity spatiotemporal sensing and analytics.

The Gridmetrics platform makes sensor data available via secure, authenticated, machine-to-machine application programming interfaces (APIs) that allow utilities and other authorized third parties to access and utilize streams of readings and anomaly events to augment their existing supervisory control and data acquisition systems and processes. Using Cable's 10G broadband network, sensor data is transported at gigabit speeds and millisecond latencies to meet the immediate need to provide real-time APIs that improve grid visibility at a time when utilities need better insights now

## 2. Fundamental Cable Architecture

Understanding the cable television network and how it works can help decision makers outside the cable industry in considering the use and application of Gridmetrics data. By way of background, cable services are enjoyed by tens of millions of U.S. households in urban, suburban, and rural areas. Cable high-speed internet is available to 90% of U.S. households and 80% of homes have access to gigabit speeds.<sup>5</sup>

Similar to the evolution of the electric power grid, cable networks started as one-way transmission paths from a central location, called a headend, to each customer's home. The current cable architecture involves extensive use of fiber optics that carry two-way interactive data with millisecond latencies. Much like the grid, the cable network topology has evolved over the years into a tree-and-branch architecture as shown in Figure 1.



**Figure 1 – Cable tree-and-branch architecture**

<sup>5</sup> <https://www.ncta.com/industry-data>


### 3. Normal Grid Operating Voltages

The American National Standards Institute (ANSI) documents nominal voltage ratings and operating tolerances for electric power systems in ANSI C84.1-2016, which is available from the National Electrical Manufacturers Association (NEMA).<sup>6</sup> As shown in Table 1, the lower and upper acceptable voltage limits are 95% and 105% respectively, and there are additional considerations for the frequency, intensity, and duration of voltage excursions.<sup>7</sup>

**Table 1 - Normal operating voltage limits [Volts AC]**

| Nominal Voltage | 95% Lower Limit | 105% Upper Limit |
|-----------------|-----------------|------------------|
| 120             | 114             | 126              |
| 240             | 228             | 252              |

Within the normal operating voltages in Table 1, the customer experience is expected to be acceptable. Above or below voltage limits, issues may arise with billing, resilience, safety, and equipment longevity. For example, high voltages lead to higher energy usage and higher electric bills and may damage capacitors on electric motors widely used in refrigeration, heating, air conditioning, and water pumps. Low voltages lead to overheating and a shortened lifespan of motors. The observed voltages in the Gridmetrics national data set regularly deviate from the NEMA boundaries and are shown in Figure 2.




**Figure 2 - Observed voltages in Gridmetrics national data set**

<sup>6</sup> <https://webstore.ansi.org/standards/nema/ansic842016>

<sup>7</sup> <https://www.spgsamerica.com/information/acceptable-voltage-ranges>

Figure 2 depicts the distribution of voltage observations. The horizontal axis represents the range of sensor voltages and the vertical axis is the number of observations of each voltage. The green vertical lines represent the target voltages of 120 VAC and 240 VAC<sup>8</sup>. In the Gridmetrics national data set, approximately 92% of sensors are connected to 120 VAC, 8% of sensors are connected to 240 VAC, and less than 1% of sensors are connected to 208 VAC.

The average voltage per sensor in the Gridmetrics national data set is shown in Figure 3.




**Figure 3 - Observed voltages in Gridmetrics national data set**

Figure 3 depicts the mean (average) voltage per sensor in the Gridmetrics national data set. All sensors are sorted by mean voltage along the horizontal axis, with the lowest voltages at the left and highest voltages at right. The vertical axis is the mean of all voltages observed by each sensor. The target values of 120 VAC and 240 VAC appear in green, and NEMA limits of  $\pm 5\%$ ,  $\pm 8\%$ , and  $\pm 10\%$  appear in red. In a perfect world, instead of voltage observations above and below the red lines, all voltages would be depicted along the flat green lines at 120 VAC or 240 VAC.

<sup>8</sup> Data used in all figures are available for select U.S. areas to allow for troubleshooting voltage quality issues and for comparison of the performance of individual utility serving areas to each other and to the national data set.

## 4. High Voltages

An example of high voltages observed by one sensor in the Gridmetrics national data set is shown in Figure 4.




**Figure 4 - Example of high voltage in the Gridmetrics national data set.**

Figure 4 depicts voltages above 240 VAC that are so high that the target voltage and some NEMA limits are not visible. The horizontal axis is one month of time expressed in 5-minute periods and the vertical axis is the observed voltage.

Utility operations procedures can include checking sensor locations for anomalies. Additional examples of locations reporting high 120 and 240 voltage ranges are collected and available.

## 5. Low Voltages

An example of low voltage observed by one sensor in the Gridmetrics national data set is shown in Figure 5.




**Figure 5 - Example of low voltage in the Gridmetrics national data set.**

Figure 5 depicts voltages below 240 VAC, so low that the target voltage and some NEMA limits are not visible. The horizontal axis is the 5-minute time period and the vertical axis is the observed voltage. Additional examples of locations reporting low 120 and 240 voltage ranges are available.

## 6. Fluctuating Voltages

Some voltages fluctuate randomly, others fluctuate periodically. An example of a periodic fluctuation observed by one sensor in the Gridmetrics national data set is shown in Figure 6.



**Figure 6 - Example of fluctuating voltage in the Gridmetrics national data set.**

In Figure 6, the 120 VAC target is denoted in green and the ANSI limits are denoted in red. The horizontal axis is time in 5-minute increments and the vertical axis is the observed voltage.

Fluctuations can be caused by loose connections, faulty appliances, and malfunctioning motors. Large fluctuations from commercial and industrial loads can affect many consumers on the same circuit. Perhaps most problematic is electric arcing from loose connections that can result in overheating, sparks, and downed wires that cause wildfires. Additional examples of locations reporting fluctuating 120 and 240 voltage ranges are available.

## 7. Outages

Interruptions in electricity service vary by frequency and duration across geographic areas and are affected by weather and other variables. For any area, outage data can be very valuable to utilities in prioritizing and justifying investments that help mitigate outage impacts as measured in their System Average Interruption Duration Index (SAIDI) and System Average Interruption Frequency Index (SAIFI). In recent years, the number of outages and the number of customers affected have consistently risen, in some areas by as much as 200%.<sup>9</sup> For example, in 2017, the average U.S. utility customer experienced 1.4 interruptions and lost their power for a total of 7.8

<sup>9</sup> <https://www.bloomenergy.com/bloom-energy-outage-map>

hours over the year, nearly double the average total duration of interruptions experienced in 2016.<sup>10</sup>

Power outages are both irritating to utility customers and regulators and are also incredibly expensive. According to a 2017 analysis by Sentient Energy, each minute of mitigated SAIDI can save a utility \$500K to \$1.5 million in O&M (operation & maintenance) costs.<sup>11,12</sup> Gridmetrics data can aid in reducing SAIDI and SAIFI by assisting utility outage prediction and detection algorithms.

A smarter grid can reduce the economic costs associated with power disturbances. Studies by the Electric Power Research Institute have estimated the yearly cost of power disturbances across all business sectors in the United States at between \$104 billion and \$164 billion as a result of outages, and another \$15 billion to \$24 billion due to power quality phenomena.

Gridmetrics data are helpful in finding and quantifying unforeseen and previously invisible outages and power quality issues in the secondary distribution grid. Gridmetrics data can be leveraged to complement existing utility outage frequency and duration processes as well as the metrics that utilities are required to report to the U.S. Energy Information Administration.

For example, Gridmetrics automated 5-minute near real-time observations during a severe storm in Sioux Falls, SD, have proven insightful. Sioux Falls was struck by three tornadoes around 11:30 PM on September 10, 2019, as depicted in Figures 7, 8, and 9.


---

<sup>10</sup> <https://www.eia.gov/todayinenergy/detail.php?id=37652>

<sup>11</sup> <https://www.sentient-energy.com/blog/15-minutes-could-save-you-15000000-or-more>

<sup>12</sup> <https://www.epri.com/research/products/1022519>






**Figure 7 - Status of the secondary distribution grid prior to the outage.**

Figure 7 shows the “green” status of sensors before the tornadoes hit. All but one sensor is green, indicating the grid power is “on” and widely available. This is a normal grid status.





**Figure 8 - Status of the secondary distribution grid during the outage.**

In Figure 8, broadband sensors show the path of the Tornadoes starting at the lower left. The location of each red dot depicts a power failure. A changing map over time reveals where power outages come and go, while clusters of red provide a sense of the size of areas affected.


A unique capability of Gridmetrics sensors is their ability to continue to report status during grid power failures. Continuous reporting by all Gridmetrics sensors during weather catastrophes and other outages is made possible by battery-backed uninterruptible power supplies in each sensor. Batteries in each sensor provide 3 to 5 hours of backup power for sensing functions as well as high-speed real-time communications for backhaul of voltage observations.

Two benefits have become apparent to users of Gridmetrics data. First, due to Gridmetrics' enhanced visibility of the secondary distribution grid, new outages are being exposed over and above those detected by and reported through existing utility energy management and SCADA (supervisory control and data acquisition) systems. As such, the additional visibility in the last mile of the grid provided by Gridmetrics is benefiting utilities and their customers.

Second, utilities generally have limited tools to verify service restoration and as such can use Gridmetrics data to automate and augment their existing processes that involve utility staff

driving around to see where power appears to have been restored. As one might expect, few utility customers call to report their service has been restored, making Gridmetrics real-time data APIs a natural fit to fill this need.

Figure 9 depicts a reduction in outages after the storm passed.



**Figure 9 - Example of service near the end of the outage**

The reduction in outages in Figure 9 is likely the result of utility service crews working to restore power around the city, especially at lower left. The remaining red dots depict serving areas that still require attention by utility crews.

## **8. Cable TV Broadband Sensor and Network Advantages**

While most utilities have some level of experience with automatic meter reading (AMR) and advanced metering infrastructure (AMI), it is important to note inherent limitations of these technologies.

A) By design, all meters used for AMR in the U.S. are limited to 240 VAC split-phase leg-to-leg measurements, providing no measurement of the performance of ground or neutral circuits. For reference, all Gridmetrics sensors are grounded, battery-backed, and provide continuous

reporting during grid outages. Over 90% of Gridmetrics sensors are connected via leg-to-neutral to 120 VAC, which enables visibility of potential loss-of-life safety issues and reduced equipment longevity issues due to insufficient or intermittent grounding and unbonded neutral circuits.

B) While more than half of U.S. households have AMR deployed, the capability is often unused in daily operations due to bandwidth limitations in the backhaul communications infrastructure that results in bottlenecks in utilities receiving and processing AMR data.<sup>13,14</sup> In contrast, all Gridmetrics sensor data are backhauled over gigabit networks with millisecond latencies, allowing for near real-time 5-minute updates for all sensors in the national data set. The latency, loss, throughput, and jitter of the Gridmetrics sensor backhaul network far outperform commonly deployed AMI networks.

## 9. Conclusion

As the number of DERs and the two-way electricity flows they create rapidly increase, aging grid infrastructure elements that are supposed to keep the grid safe are failing and causing unprecedented loss of life and property.<sup>15</sup> While the enormity of the electric power grid is such that in the U.S. alone, the 5.5 million mile distribution network is long enough to reach the moon nearly 21 times--the performance of the last mile of the grid is sparsely monitored and hence unable to be optimally managed. Gridmetrics 5-minute sensor readings fill the immediate need to augment utility supervisory control and data acquisition systems by rapidly improving the monitoring of the secondary distribution portion of the grid by using an existing fleet of power sensors deployed across broadband networks.

The growing and evolving Gridmetrics data set is available to aid U.S. utilities in monitoring and managing the secondary distribution networks that make up the last mile of the grid. The locations of specific anomalies worthy of investigation are available for use in utility operations. Through utility maintenance and repair efforts, infrastructure aging, wear and tear, and local weather--the location and severity of anomalies will change over time, supporting the case for real-time Gridmetrics APIs and near real-time data feeds. Through collaboration with utilities and sharing best practices for anomaly detection and classification, it is expected that anomalies that foretell of impending infrastructure failures, resiliency and safety issues, and high-risk for loss-of-life can be identified. In addition, working with utilities, the criteria used to identify anomalies can be expanded, refined, and validated to achieve maximum benefit from Gridmetrics data.

Data from broadband power quality sensors can help utilities pinpoint existing portions of the grid that can be inspected for high, low, and fluctuating voltages which can cause unsafe conditions, poor customer experiences, and premature failures of customer equipment. In addition, outage data from broadband sensors can be correlated with existing utility data sets to create a more comprehensive understanding of distribution network frailties. Combining insights

---

<sup>13</sup> <https://www.smart-energy.com/top-stories/ami-system-operations-pose-unexpected-challenges/>

<sup>14</sup> <https://pubs.naruc.org/pub/FA865F09-934E-F89E-5141-7E766D260068>

<sup>15</sup> <https://www.tdworl.com/distributed-energy-resources/article/21120102/democratizing-energy-the-rise-of-ders>

from utility SCADA systems and Gridmetrics data can help improve network reliability, resilience, and safety.

## Abbreviations

|       |                                               |
|-------|-----------------------------------------------|
| AMI   | advanced metering infrastructure              |
| AMR   | automatic meter reading                       |
| ANSI  | American National Standards Institute         |
| API   | application programming interface             |
| DER   | distributed energy resource                   |
| EV    | electric vehicle                              |
| O&M   | operations and maintenance                    |
| SAIDI | System Average Interruption Duration Index    |
| SAIFI | System Average Interruption Frequency Index.  |
| SCADA | Supervisory Control and Data Acquisition      |
| SCTE  | Society of Cable Telecommunications Engineers |
| U.S.  | United States                                 |
| VAC   | volts alternating current                     |

## Bibliography & References

*U.S. Electrical Grid Undergoes Massive Transition to Connect to Renewables*, Jennifer Weeks, Scientific American, The Daily Climate, April 28, 2010.  
<https://www.scientificamerican.com/article/what-is-the-smart-grid/>

*PG&E Knew for Years Its Lines Could Spark Wildfires, and Didn't Fix Them*, Katherine Blunt and Russell Gold, The Wall Street Journal, July 10, 2019. <https://www.wsj.com/articles/pg-e-knew-for-years-its-lines-could-spark-wildfires-and-didnt-fix-them-11562768885>

*This Old Metal Hook Could Determine Whether PG&E Committed a Crime*, Russell Gold and Katherine Blunt, The Wall Street Journal, March 8, 2020. <https://www.wsj.com/articles/this-old-metal-hook-could-determine-whether-pg-e-committed-a-crime-11583623059>

ANSI C84.1-2016: *Electric Power Systems And Equipment - Voltage Ratings (60 Hz)*; NEMA: National Electrical Manufacturers Association. Accessed July 16, 2020.  
<https://webstore.ansi.org/standards/nema/ansic842016>

*Acceptable Voltage Ranges*, Safety Protection Grid Solutions. Accessed July 16, 2020.  
<https://www.spgsamerica.com/information/acceptable-voltage-ranges>

*California Power Outage Map*, Bloom Energy. Accessed July 16, 2020.  
<https://www.bloomenergy.com/bloom-energy-outage-map>

*Average U.S. electricity customer interruptions totaled nearly 8 hours in 2017*, Today in Energy, U.S. Energy Information Administration, November 30, 2018.  
<https://www.eia.gov/todayinenergy/detail.php?id=37652>

*15 Minutes Could Save You \$15,000,000 or More.....*, Sentient Energy Blog, August 22, 2017.  
<https://www.sentient-energy.com/blog/15-minutes-could-save-you-15000000-or-more>

*Estimating the Costs and Benefits of the Smart Grid: A Preliminary Estimate of the Investment Requirements and the Resultant Benefits of a Fully Functioning Smart Grid*, Electric Power Research Institute, 2011. <https://www.epri.com/research/products/1022519>

*AMI system operations pose unexpected challenges*, David Gordon Kreiss and Masoud Abaei, Smart Energy International, September 23, 2013. <https://www.smart-energy.com/top-stories/ami-system-operations-pose-unexpected-challenges/>

*Toward an End-to-End Smart Grid: Overcoming Bottlenecks to Facilitate Competition and Innovation in Smart Grids*, Johann J. Kranz and Arnold Picot, National Regulatory Research Institute, June 2011. <https://pubs.naruc.org/pub/FA865F09-934E-F89E-5141-7E766D260068>

*Democratizing Energy: The Rise of DERs*, Stephan Marty, T&D World, January 9, 2020.  
<https://www.tdworld.com/distributed-energy-resources/article/21120102/democratizing-energy-the-rise-of-ders>

# **Powering the Near Future 10G Access Network**

## **Considerations for Assuring Sufficient and Reliable Power**

A Technical Paper prepared for SCTE•ISBE by

**Rob Anderson**

Director of Product Management  
EnerSys Energy Systems  
3767 Alpha Way, Bellingham, WA 98226  
360-392-2293  
randerson@alpha.com

# Table of Contents

| Title                                                                                   | Page Number |
|-----------------------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                                    | 4           |
| 2. The Near Future 10G Network .....                                                    | 4           |
| 2.1. Distributed Access Architecture (DAA) and Distributed CCAP Architecture (DCA)..... | 5           |
| 2.2. Remote PHY and Flexible MAC Architecture .....                                     | 5           |
| 2.3. Coherent Optics Backhaul .....                                                     | 6           |
| 2.4. Generic Access Platform (GAP) Nodes.....                                           | 7           |
| 2.4.1. RAN Module .....                                                                 | 8           |
| 2.4.2. EPON Module .....                                                                | 8           |
| 2.4.3. Edge Compute Module.....                                                         | 8           |
| 2.5. DAA Node Power Budget Summary.....                                                 | 9           |
| 3. Powering The 10G Access Network .....                                                | 9           |
| 3.1. Traditional HFC Power Review .....                                                 | 9           |
| 3.1.1. HFC Power Supplies.....                                                          | 10          |
| 3.1.2. HFC Powering Example .....                                                       | 11          |
| 3.2. Powering the Evolving Access Network .....                                         | 13          |
| 3.2.1. Node Splits .....                                                                | 13          |
| 3.2.2. Deep Fiber (N+0).....                                                            | 13          |
| 3.2.3. 1.8 GHz Extended Spectrum.....                                                   | 16          |
| 3.2.4. 5G Fixed Wireless Access.....                                                    | 16          |
| 3.2.5. EPON .....                                                                       | 18          |
| 4. Ensuring Reliable Power .....                                                        | 19          |
| 4.1. Power Reliability Considerations.....                                              | 19          |
| 4.1.1. Intra-Node Power Hold-Up Time .....                                              | 19          |
| 4.1.2. Utility Backup Time .....                                                        | 19          |
| 4.1.3. Redundant Power Source .....                                                     | 20          |
| 4.2. Backup Power Requires Healthy Batteries.....                                       | 20          |
| 4.2.1. Effects of Charge on Battery Capacity.....                                       | 20          |
| 4.2.2. Battery Chemistry and Charge Mismatch.....                                       | 21          |
| 4.2.3. Preventative Maintenance is Essential .....                                      | 23          |
| 5. Conclusion.....                                                                      | 24          |
| Abbreviations.....                                                                      | 26          |
| Bibliography & References .....                                                         | 27          |

## List of Figures

| Title                                                | Page Number |
|------------------------------------------------------|-------------|
| Figure 1 – 10G Influenced Technologies.....          | 5           |
| Figure 2 – DAA Configurations .....                  | 6           |
| Figure 3 – ODC/Virtual Hub Architecture.....         | 7           |
| Figure 4 – Edge Compute Concept from CableLabs ..... | 8           |
| Figure 5 – Traditional HFC Access Network.....       | 9           |
| Figure 6 – Typical Pole Mounted Broadband UPS.....   | 10          |
| Figure 7 – HFC Powering Example .....                | 11          |
| Figure 8 – DAA N+0 Powering Example .....            | 14          |
| Figure 9 – Coax Power Bridging .....                 | 16          |

|                                                     |    |
|-----------------------------------------------------|----|
| Figure 10 – 5G Fixed Wireless Access Powering ..... | 17 |
| Figure 11 – EPON Block Diagram.....                 | 18 |
| Figure 12 – Charge Effect on Battery Life .....     | 21 |
| Figure 13 – Battery Charge Configuration .....      | 22 |
| Figure 14 – Internal Battery Resistance .....       | 22 |
| Figure 15 – Battery Corrosion .....                 | 23 |

## List of Tables

| <b>Title</b>                                              | <b>Page Number</b> |
|-----------------------------------------------------------|--------------------|
| Table 1 – DAA Node Power Budget.....                      | 9                  |
| Table 2 – Coax Spans for HFC Powering Example .....       | 12                 |
| Table 3 – HFC Powering Example Results.....               | 13                 |
| Table 4 – Coax Spans for DAA N+0 Powering Example.....    | 14                 |
| Table 5 – DAA N+0 Powering Example Results .....          | 15                 |
| Table 6 – Modified DAA N+0 Powering Example Results ..... | 15                 |
| Table 7 – 5G Fixed Access Wireless Powering Results.....  | 18                 |



## 1. Introduction

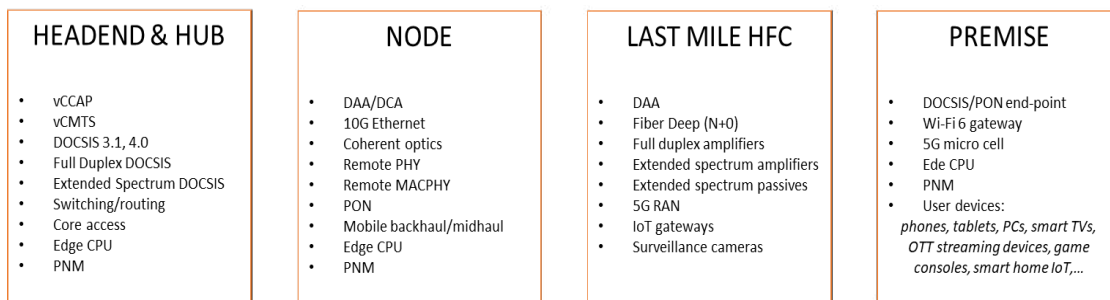
The 10G initiative is the catalyst behind several technology innovations designed to deliver future proof internet speeds up to 100 times faster than most consumers are experiencing today. These innovations will affect every aspect of the broadband network including headends, the access network and the customer premises. The access network specifically must undergo enhancements to support performance levels envisioned by the 10G initiative. Many of these network revisions involve technology that is either new or under development. Operators will need to make decisions about implementing new technology to be certain that their networks will be able to handle whatever is coming. A service provider's existing network architecture along with the cost for upgrades will steer operators towards the technologies right for their needs. Network upgrade options may include anything from standard node splits and remote PHY (R-PHY) node upgrades to potentially disruptive technologies like 5G fixed wireless or fiber overlays. Regardless of the network upgrade specifics, power is a common requirement for any network.

Assuring the availability of additional, reliable and intelligent power for the 10G capable network is both essential and challenging since network architectures are evolving and much of the 10G enabling technology is still being developed. In this paper we address the access network powering challenge by providing operators with a set of powering guidelines. Our objective is to help ensure that sufficient, reliable network power is available, irrespective of the specific technologies implemented to meet network performance objectives. We accomplish this by first describing a 10G reference access network. Our reference network includes architectures and technologies that are both current and that are under development. Next, we overlay our reference network with the appropriate powering architecture. Existing hybrid fiber coax (HFC) powering infrastructure is re-used wherever practical. Some new 10G network elements may require new and innovative powering options. While no single service provider would incorporate every element of our 10G reference network, operators can plan powering strategies to support those network upgrades that will meet their future performance goals.

## 2. The Near Future 10G Network

In early 2019 the 10G initiative was announced with the objective of providing 10Gbps symmetrical, secure, low latency data services. CableLabs coined the phrase near future in connection with 10G. A CableLabs paper provides the following description of the 10G near future network: "Immersive experiences like interactive holographic projections, video walls and next-generation artificial intelligence (AI) and virtual reality (VR) tools will all require a super high-capacity network that can deliver an immense amount of data to the end user ... 10G technologies will enable speeds 10 times faster than the 1 gigabit speeds that cable offers today, equally applied to both the upstream and downstream traffic over existing cable hybrid fiber coax (HFC) networks [1]."

Our 10G reference access network includes existing technologies and technologies that are under development. Figure 1 – 10G Influenced Technologies, highlights some 10G influenced technologies and their relative location in the broadband network.



**Figure 1 – 10G Influenced Technologies**

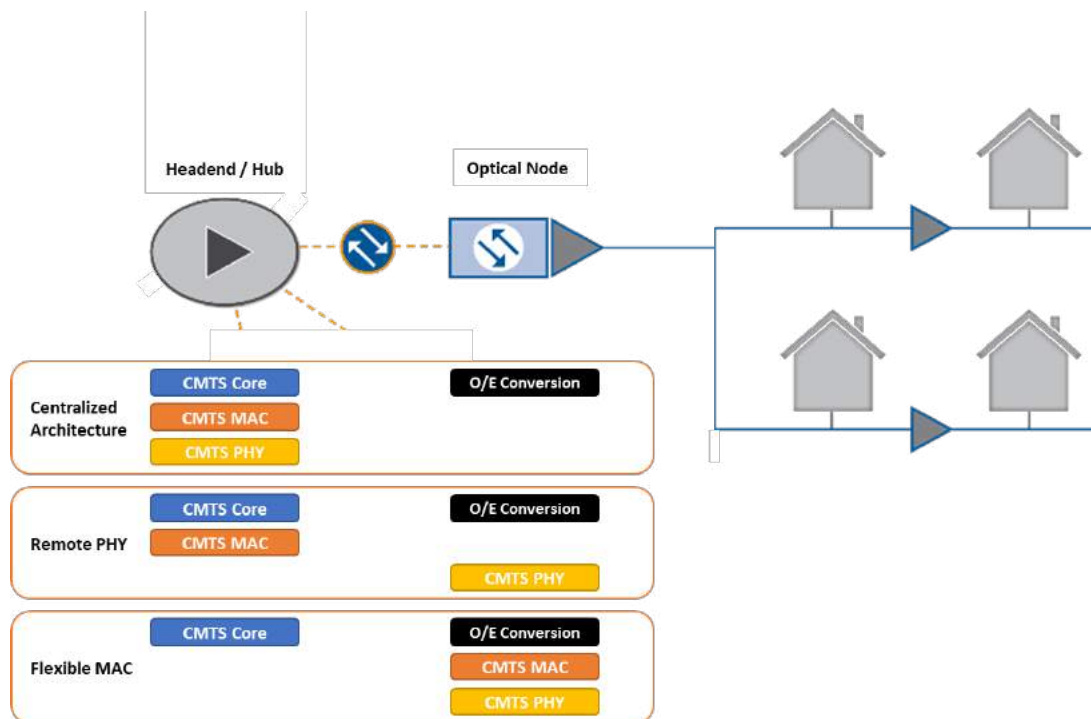
Our discussion will focus on elements contained in the node and last mile HFC blocks shown in Figure 1. These access network elements are within the scope of our powering discussion.

## **2.1. Distributed Access Architecture (DAA) and Distributed CCAP Architecture (DCA)**

Distributed Access Architecture (DAA) relocates or distributes functions that traditionally reside in the headend or hub closer to the user. Moving functions deeper into the network reduces power and space demands in the headend/hub. As functions move closer to users the network experiences increase in efficiencies, speed, reliability, latency and security [2]. Distributed CCAP Architecture (DCA) outlines specific technologies to distribute headend functions to the distributed network [3]. For the intent of our powering discussion DAA and DCA will be considered the same and we shall use the term DAA going forward.

## **2.2. Remote PHY and Flexible MAC Architecture**

Remote PHY (R-PHY) and Flexible MAC Architecture (FMA) are implementations under the DAA umbrella. Both technologies move processing functions from the headend/hub to distributed optical nodes as shown in this diagram.



**Figure 2 – DAA Configurations**

Operators transitioning from traditional analog nodes to DAA nodes must increase their node power budget. Analog nodes consuming about 80W each are replaced by DAA nodes requiring more power (assuming similar configurations). Current industry data indicates that DAA node power consumption is in the range of 140W to 190W with FMA capable nodes requiring more power than R-PHY capable nodes. As field programmable gate arrays (FPGA) and application specific integrated circuits (ASIC) evolve, this power consumption is expected to improve. For our DAA node power budget we'll assume a nominal value of 165W per node.

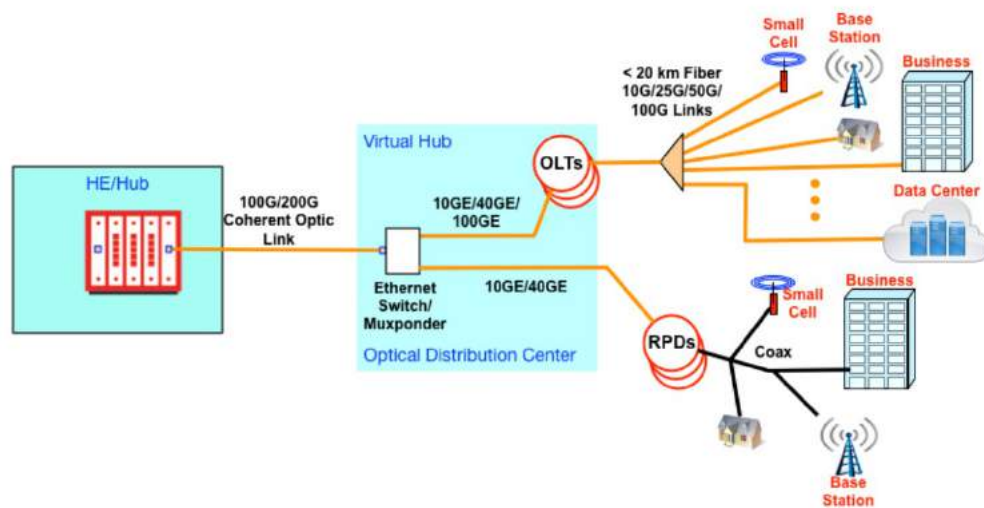
At the time of this paper's publication DOCSIS™ 4.0 has not yet been implemented in DAA nodes. DOCSIS 4.0 radio frequency (RF) signaling, including Full Duplex DOCSIS (FDD) and Extended Spectrum DOCSIS (ESD) with downstream spectrum support to 1,794 MHz, are generated in the PHY component of the DAA node [4]. Specific power requirements for DOCSIS 4.0 enabled DAA nodes are not known. For estimating purposes we'll assume that future DOCSIS 4.0 enabled DAA node power requirements fit within our estimated DAA power budget of 165W.

Beyond the current DOCSIS 4.0 standards, there are industry proposals to extend the downstream spectrum beyond 1.8GHz to 3GHz. This next-generation RF band upgrade could enable upstream speeds in excess of 9 Gbps and downstream speeds to 25 Gbps. The proposal utilizes small RF amplifiers near the premises, providing signal amplification for the 1.2GHz to 3GHz range [5]. The impact of this enhanced RF spectrum on the HFC power network would be speculative at this stage so we will leave powering this concept for a future discussion.

### 2.3. Coherent Optics Backhaul

Coherent optics, or full duplex point to point (P2P) coherent optics, uses amplitude and phase to carry large amounts of bidirectional data over a single fiber using a single wavelength. In the DAA cable environment, coherent optics will be used to establish high-capacity links from the headend/hub to an

aggregation node or directly to the DAA node. The current state of the industry supports 100Gbps outdoor rated coherent optics pluggable modules with 200Gbps and 400Gbps modules planned for the future. Our access network powering model will consider two potential coherent optics use cases. First is a coherent optics module providing high speed backhaul direct to the DAA. Today's coherent optics modules consume approximately 15W which will be added to our DAA node power budget. Next, we consider the addition of an optical distribution center (ODC) or virtual hub containing the coherent optical components as well as Ethernet switching elements. An access network with an ODC is pictured here [6].



**Figure 3 – ODC/Virtual Hub Architecture**

In this network the ODC is shown supporting optical outputs to DAA nodes (RPDs) as well as supporting OLT functions for passive optical network (PON) network segments. For power budgeting purposes the ODC will be considered a strand mount style element with power consumption of 150W. Powering the ODC or other virtual hubs can be more challenging than powering optical nodes. With the ODC, there are no coax connections to source power to any of the fiberoptic inputs and outputs. If the ODC is installed in a location formerly occupied by an optical node then powered coax may be available at the location to power the ODC. If powered coax is not available, alternative powering options can include anything from installing a dedicated power cable from nearby powered coax to installing a new uninterruptable power supply (UPS) at the ODC location. Irrespective of the power input source, the operator will consider the ODC a critical location requiring extremely reliable power due to the number of downstream subscribers serviced by this device.

## 2.4. Generic Access Platform (GAP) Nodes

The SCTE Interface Practices Subcommittee, Working Group 1 (IPS WG1) is defining a set of specifications for a Generic Access Platform (GAP) node housing. One objective of the GAP node is “increased availability and ability to integrate advanced technologies within a modular approach [7].” Integrating additional technologies inside the node housing will increase power consumption which is why the GAP is mentioned here. Future plug-in modules discussed during GAP committee sessions include: radio access network (RAN) (Wi-Fi, 5G and citizen broadband radio service (CBRS)), Ethernet passive optical network (EPON) and edge computing modules. These three potential node enhancements are discussed here.

### 2.4.1. RAN Module

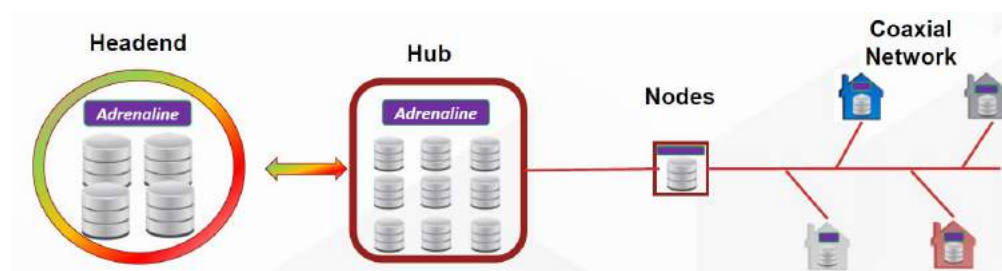
A potential access network upgrade includes a RAN component, potentially taking the form of 5G fixed wireless radio units (RU's) utilizing the coax for power and backhaul. It is not likely the DAA node physical location will align with RU location requirements for signal coverage. For simplicity we assume the DAA does not incorporate wireless modules. A 5G fixed wireless access network extension is discussed in section 3.4.2.

### 2.4.2. EPON Module

EPON modules could be installed in a DAA (GAP or traditional) node enclosure. A challenge comes with attempting to define a standard configuration. Some EPON components are light weight, layer 2 devices that combine with cloud based processing for full featured performance. One module from a well-known component vendor supports 2x10GigE x 2x10GEAPON and consumes only 35W. Specifications for more full featured EPON units supporting 4x10GigE x 4x10GEAPON and including native processing support for features such as DOCSIS provisioning over Ethernet (DPoE) and layer 3 networking have been identified as consuming up to 140W. To avoid numerous permutations of EPON and RF QAM node configurations and to simplify our powering discussion it is assumed that our DAA node may support one basic EPON module at 35W.

### 2.4.3. Edge Compute Module

CableLabs in 2020 announced a program called Adrenaline intended to “transform the network into a distributed heterogeneous compute platform with dynamic workload allocation”. The Adrenaline concept is shown here [8].



**Figure 4 – Edge Compute Concept from CableLabs**

An Adrenaline type distributed, or edge computing concept has been the topic of multiple industry webinars. Edge computing is distributed to intelligent components throughout the network, including the DAA node. The cable broadband industry is a few years away from realizing speed and performance improvements from distributed edge computing, but many are working towards this goal. For our access network powering budget, we assume our near future DAA node will be fitted with a processing module. Power estimates at this stage are speculative but let's assume our DAA processing element will add an additional 10W power requirement to the node.

## 2.5. DAA Node Power Budget Summary

From our near future access network concept node review the following table summarizes our DAA node powering budget.

**Table 1 – DAA Node Power Budget**


| Configuration                  | Power (W) |
|--------------------------------|-----------|
| DAA (RPHY/ FMA) node           | 165       |
| Coherent optical node backhaul | 15        |
| Edge compute module            | 10        |
| Total DAA node power budget    | 190       |
| Adding EPON                    | 35        |
| DAA with RF and EPON           | 225       |

Our DAA node of the future requires significantly more power than traditional analog nodes. This difference combined with other related access network upgrades give cause to review powering assumptions and consider evolving needs.

## 3. Powering The 10G Access Network

### 3.1. Traditional HFC Power Review

The access network includes all network elements beyond the headend/hub and up to the premises. Some network architectures utilize virtual hubs housed within outdoor sealed enclosures which are also considered part of the access network. A simplified access network block diagram is shown in Figure 5.



**Figure 5 – Traditional HFC Access Network**

In the traditional HFC access network signals from the headend/hub are transmitted over fiber optic cable to optical nodes. Nodes convert between optical and electrical signals and convey those signals to multiple coax network segments. Each coax segment typically consists of a tree-and-branch structure for signal distribution. Passive splitters subdivide the signal into multiple paths (branches). Amplifiers are located throughout the coax network segment to boost signals, providing appropriate signal levels to the end users, which are typically homes and businesses. Directional taps located near the end users will “tap off” the signal from the main coax cable into drop cables which bring the signals into the customer premises. Nodes and amplifiers require power. Power is also consumed by coax line loss from Joule heating ( $I^2R$  losses). Power and RF signals are both multiplexed onto the coax, eliminating the need for separate power and signal cables. Power supplies are placed as needed throughout the access network to provide power to nodes and amplifiers.

### **3.1.1. HFC Power Supplies**

Power for HFC components is provided by a specialized backup power system known as a broadband UPS. UPS systems are physically located throughout the coax portion of the network where required to provide power for each active network element. UPSs are physically installed on outdoor utility poles, in dedicated ground mounted enclosures, and in secured utility areas of multi-dwelling units (MDUs). A typical utility pole mounted broadband UPS is shown in Figure 6.



**Figure 6 – Typical Pole Mounted Broadband UPS**

The UPS converts utility power (120VAC or 240VAC in North America) to 90VAC for insertion into the coax. Early cable networks used 30VAC then later 60VAC for network power. Today, networks almost exclusively use 90VAC power with a few older 60VAC networks still in operation.

Unlike utility service which provides sinusoidal AC power, the broadband UPS produces a quasi-square wave or trapezoidal shaped power output. This wave shape is a result of the ferro-resonant transformer used in broadband UPS systems. The ferro-resonant transformer provides a high level of electrical isolation, protecting powered nodes and amplifiers from utility line power surges and transients that could damage sensitive electronics.

Batteries within the UPS system enable the power supply to provide continuous, reliable backup power to the access network during utility disruptions. The output of the UPS is connected to a power inserter, which acts like a reverse directional tap, to inject power onto the coax cable. Some operators use the term “shunt” to describe the process of injecting power or shunting power onto the coax.




Each amplifier, splitter and tap can be configured to pass power through itself and on to the next device in the network or to block power from passing through itself. The decision to pass or block power within a specific network device is determined by the operator and is based on criteria including:

- Do downstream devices require power and if so, how much?
- Does the UPS have sufficient capacity to power future planned devices?
- Will powering additional devices cause this UPS to exceed the cable broadband operator's maximum powering policy?

### 3.1.2. HFC Powering Example

Figure 7 uses a simplified HFC network segment to illustrate some basic powering concepts. In practice, HFC powering is more complex than illustrated but several basic principles are shown.



**Figure 7 – HFC Powering Example**

As stated, the diagram represents a simplified HFC network segment. In this example a broadband UPS is connected near the optical node. The node is powered from this coax segment. Amplifiers 1-4 are also powered from this UPS. Amplifier 1 and amplifier 2 are each configured to pass power through their chassis, enabling downstream devices to be powered from the UPS. The splitter is configured to pass power through both outputs. Amplifiers 3 and 4 are powered from the UPS and are configured to block power from passing to their respective outputs eliminating additional power draw from any components further down the coax.

Assume that the node requires 80W of power and amplifiers 1-4 each require 70W of power to operate. Also assume the UPS output is configured to 90VAC and that it has the capacity to provide up to 1350W of power. For simplicity, ignore coax line loss. The total power required is calculated as the sum of power required from each network active:

$$P(\text{Actives}) = P(\text{node}) + P(\text{amp1}) + P(\text{amp2}) + P(\text{amp3}) + P(\text{amp4})$$

$$\text{Total Power (Actives)} = 80\text{W} + (70\text{W} \times 4) = 360\text{W}$$



To make our example more realistic we include coax line loss in our power equation. Assume all coax is 0.625 inches diameter which has a typical resistance of 0.0011 ohms/feet. Also, assume the following coax span lengths.

**Table 2 – Coax Spans for HFC Powering Example**

| From     | To       | Span (ft) |
|----------|----------|-----------|
| UPS      | node     | ~0        |
| node     | amp 1    | 1,000     |
| amp 1    | amp 2    | 1,000     |
| amp 2    | splitter | 1,000     |
| splitter | amp 3    | 1,000     |
| splitter | amp 4    | 1,000     |

Note: Using these assumptions, the total coax length between the UPS and amplifier 3 or the UPS and amplifier 4 is 4000ft.

Let's calculate the power lost in a single 1,000ft segment of coax. Using Ohm's Law:

$$P(\text{loss}) = I^2R$$

Where:

$P(\text{loss})$  = power lost from coax line resistance, measured in Watts. Note: this energy is converted to heat, hence the term Joule heating for  $I^2R$  losses.

$I$  = current through the cable, measured in amps

$R$  = resistance of the length of cable, measured in Ohms

Taking amplifier 1 in isolation, we calculate the power lost in the coax segment between the UPS and amplifier 1 as follows:

$$P(\text{loss}) = (70\text{W}/90\text{V})^2 \times (0.0011 \text{ ohm/ft} \times 1,000\text{ft}) = 0.67 \text{ W}$$

In this example we used the Ohm's law relationship:  $I = P/V$  for the first term.

$I=P/V$  is an ideal approximation. In real-world calculations we must account for accumulative voltage drops across each coax segment, i.e., the 90VAC at the UPS output is reduced through each coax segment. The voltage drop is proportional to both cable resistance and current per the relationship:

$$V(\text{drop}) = I(\text{cable}) \times R(\text{cable})$$

HFC active elements, including nodes and amplifiers, are typically constant power devices. As input voltage to a device is reduced due to coax line resistance, the device's current will increase to maintain the required power load ( $P=VI$ ). As the current increases, power loss through the coax increases. Recall this relationship discussed earlier:  $P(\text{loss}) = I^2R$ .

In this example the node and four amplifiers require 360W to operate. Additional energy consumed (lost) due to voltage drops across the various coax segments can be calculated to be an additional 36W. Detailed calculations have been omitted for brevity. Results are summarized as in Table 3.

**Table 3 – HFC Powering Example Results**

| <b>Configuration</b>       | <b>PS I(out)</b> | <b>EOL Voltage</b> | <b>Actives Load</b> | <b>I<sup>2</sup>R Loss</b> | <b>PS Utilization</b> |
|----------------------------|------------------|--------------------|---------------------|----------------------------|-----------------------|
| Analog node, 4x Amplifiers | 4.4A             | 79.4V              | 360W                | 35.5W                      | 29%                   |

With the UPS capacity of 1350W we've consumed 29% of the available power. End of line (EOL) voltage is another important parameter to monitor. EOL voltage is the input voltage of the last active device in the network and must remain above 45V for most HFC equipment to operate. For this example, we're well within acceptable operating parameters for UPS power and EOL voltage. This approach to HFC powering has worked well for many years. However, as the access network is upgraded to support our future 10G performance goals we will soon see how our powering approach must also be revised.

### **3.2. Powering the Evolving Access Network**

Traditional HFC powering assumptions and methods must be reviewed as the underlying technology and in some cases the topology of the access network evolves to support higher performance and additional services. Section 3.2 examines a few of those network upgrades and their effect on the underlying power.

#### **3.2.1. Node Splits**

Node splits are familiar to most operators as a method of increasing upstream capacity by reducing a node service group size. If one node services 400 homes, then two nodes can service 200 homes each. Benefits come from reducing upstream (US) traffic contention. US DOCSIS utilizes time division multiple access techniques to enable many users to share the same US frequencies. Each user is granted specific time slots for US transmissions. Smaller service groups through node splits mean fewer users contending for shared upstream capacity.


Powering new nodes after a node split is usually straightforward. Existing broadband UPS systems are often already installed within existing node segments. New nodes often get installed closer to the premises and may result in eliminating amplifiers which can offset incremental power required for the new nodes.

Node splits have limitations. US bandwidth is capped by the US duplex frequency defined in each network. A low split 5-42MHz US diplexer can support about 100Mb/s of data while a 5-85Mhz mid-split diplexer can handle roughly 200Mb/s of US data. Multi-gigabit US performance requires a different approach.

#### **3.2.2. Deep Fiber (N+0)**

Expanding the node split concept to the point where all network amplifiers have been replaced by optical nodes leads to a deep fiber architecture consisting of a node + zero amplifier (N+0) configuration. Using DAA nodes in this configuration results in powerful, high performing nodes closer to the end user than ever before. With N+0 there are only passive components and coax between the DAA node and the premises. Potential data speeds are substantially higher than those possible from the traditional HFC network. Using technologies such as full duplex DOCSIS (FDD), extended spectrum DOCSIS (ESD) or both together can bring network performance much closer to the vision of the 10G concept.

However, powering an N+0 network may not be a straightforward transition from our traditional HFC powered network. Let's review a simple example to illustrate some potential challenges.



**Figure 8 – DAA N+0 Powering Example**

For our N+0 powering example we use the future DAA node defined in Table 1. Also, we're reusing the HFC node + amplifier layout from our HFC powering example but replacing the amplifiers with DAA nodes. Figure 8 shows the coax path to each DAA node. The coax is used for power only. Separate fiber connections not shown in this example carry the high speed data signals to each node.

Assume the UPS capacity is 1,350W and is configured for 90V output. Also assume that the EOL voltage must remain >45V for the nodes to operate. With our future DAA node requiring 190W, power for the active network elements is calculated as follows:

$$P(\text{Actives}) = P(\text{node1}) + P(\text{node 2}) + P(\text{node 3}) + P(\text{node 4}) + P(\text{node 5})$$

$$\text{Total Power (Actives)} = 190\text{W} \times 5 = 950\text{W}$$

To include coax line loss in our power equation, assume all coax is 0.625in diameter with resistance of 0.0011 ohm/ft. Assume the following coax span length:

**Table 4 – Coax Spans for DAA N+0 Powering Example**

| From     | To       | Span (ft) |
|----------|----------|-----------|
| UPS      | Node 1   | ~0        |
| Node 1   | Node 2   | 1,000     |
| Node 2   | Node 3   | 1,000     |
| Node 3   | Splitter | 1,000     |
| Splitter | Node 4   | 1,000     |
| Splitter | Node 5   | 1,000     |

Note: Using these assumptions, the total coax length between the UPS and Node 4 or Node 5 is 4,000ft.

In this example the total power needed for the DAA nodes is 950W. However, higher current flow through the coax results in 560W of  $I^2R$  power loss. Total power for this scenario is 1,510W which exceeds our UPS' capacity making this configuration invalid. Results are summarized here:

**Table 5 – DAA N+0 Powering Example Results**

| <b>Configuration</b> | <b>PS I(out)</b> | <b>EOL Voltage</b> | <b>Actives Load</b> | <b>I<sup>2</sup>R Loss</b> | <b>PS Utilization</b> |
|----------------------|------------------|--------------------|---------------------|----------------------------|-----------------------|
| DAA node x 5, N+0    | 17A              | 50.1V              | 950W                | 560W                       | 133%                  |

In this example the 1,350W UPS would be overloaded to 133% of capacity. It's also noteworthy that the EOL voltage has dropped to 50.1V. This is still above the 45V minimum but should be monitored as alternative powering is considered.

One option to solve our N+0 power overload is to install a second UPS and distribute the load between multiple power supplies. Adding UPS systems to the network is usually the operators last resort. Unplanned equipment costs and installation delays due to permits and local regulations can be problematic. Rather than default to installing new UPS systems as our corrective action let's look at another approach. By removing one of the DAA nodes from our test example we can solve part of our problem. Let's remove node 5 from our example, leaving the UPS to power nodes 1-4. Results follow:

**Table 6 – Modified DAA N+0 Powering Example Results**

| <b>Configuration</b> | <b>PS I(out)</b> | <b>EOL Voltage</b> | <b>Actives Load</b> | <b>I<sup>2</sup>R Loss</b> | <b>PS Utilization</b> |
|----------------------|------------------|--------------------|---------------------|----------------------------|-----------------------|
| DAA node x 4, N+0    | 10.5A            | 68.5V              | 760W                | 171W                       | 70%                   |

By removing one node our powering model now works and results in our UPS loaded to 70% capacity. Many operators place design limits on UPS utilization of between 80-85% of the UPS' capacity to reserve overhead for unplanned changes.

So how do we power the stranded node from our example? Keep in mind that this example was created specifically to illustrate the effects of I<sup>2</sup>R losses. In a live network many variables could affect the results. A few possible outcomes include:


- As the N+0 node architecture was implemented there would likely be multiple amplifiers decommissioned as part of the network upgrade. It's probable that the UPS would have spare capacity to power new nodes.
- Surplus UPS power from an adjacent coax segment could be routed to the new node(s).
- The UPS could be upgraded to a higher capacity to support the load and the required overhead (spare) power reserves.

Some operators have incorporated powering design rules into their network upgrade procedures to address these types of powering concerns. One major North American operator has established these options to address powering issues during advanced node upgrades:

- All power designs will assume advanced nodes will use existing power supplies. If any powering issues arise, the following will be allowed in order of preference:
  - Repowering of the nodes from adjacent node/power supply areas
  - 0.875" coax or dedicated power cable may be placed to reduce I<sup>2</sup>R losses
  - 15-Amp power supplies may be changed to 18-Amp at existing locations fed by 120VAC
  - 15 or 18-Amp power supplies may be added to the network
  - 15-Amp power supplies may be changed to 24-Amp at existing locations fed by 240VAC
- To prevent an overcurrent of line passives or active devices, all 24-Amp power supplies must include a protective interface module (PIM) to split current output and provide two

programmable, independent power outputs providing isolation between outputs so that a short circuit or other power anomaly in one section of the HFC plant will not cause a disruption of service to all plant fed from that power supply.

If our sample unpowered node were in an actual network, it's probable that a nearby coax segment would have enough power capacity to handle one additional node. Power from an adjacent coax segment can be bridged to the unpowered node using a technique similar to Figure 9:



**Figure 9 – Coax Power Bridging**

This power bridging technique was developed by a major North American cable operator to solve the type of power shortage problem outlined in our example. Using this method power is bridged from one distribution leg to another while RF signals are blocked. RF signals from the node connected to feeder line 1 will support customers downline from the left tap while RF signals from the node connected to feeder line 2 support customers downline from the right side tap.

### **3.2.3. 1.8 GHz Extended Spectrum**

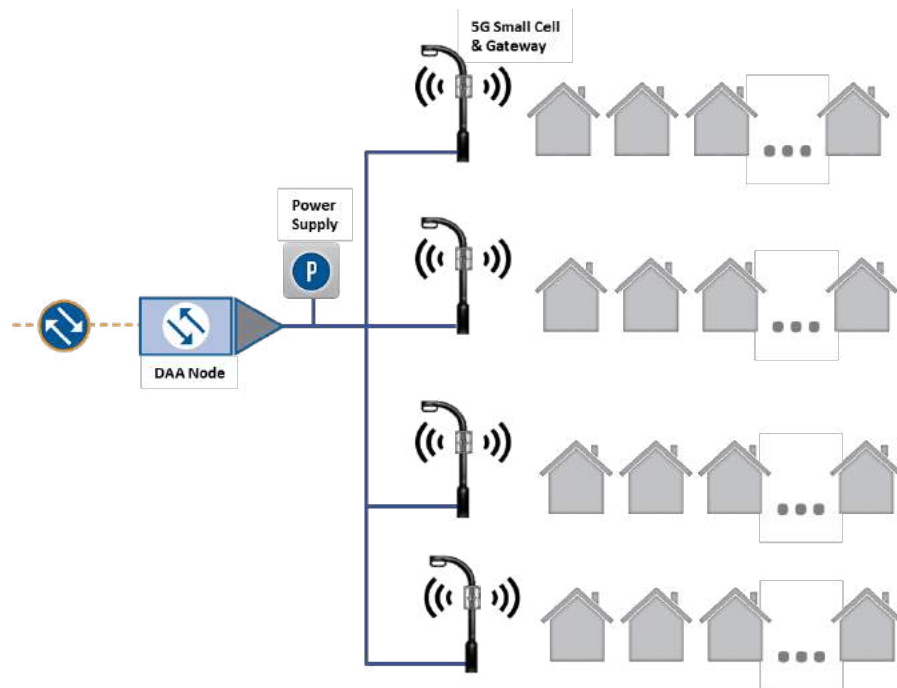
A widely discussed approach for increasing access network capacity is extending the upper limit of the usable RF spectrum to 1,794 MHz (1.8GHz). This approach is well defined in the DOCSIS 4.0 specifications [4]. What is not yet well understood at the time of this writing is the access network topology that will support the 1.8GHz extended spectrum. The DOCSIS specifications allow for different transmission profiles for frequencies above 1,002 MHz. Amplifiers capable of supporting the extended spectrum may be configured in such a way as to support higher frequencies at limited power levels. Once extended spectrum equipment is available, operators can evaluate network specific implementations to achieve their objectives. Then, the effect of extended spectrum implementations on power can be determined.

### **3.2.4. 5G Fixed Wireless Access**

The final communications link connecting end user devices such as phones, tablets and laptops is wireless. Wi-Fi 6 boasts peak data speeds in excess of 9 Gbps. Some operators are planning to deploy in-home 5G femtocells. 5G peak data rates can exceed 10 Gbps. A 5G small cell serving multiple homes on a street or throughout a neighborhood could realize similar data rates. This approach is called fixed wireless access (FWA). FWA is the process of providing wireless broadband using an RF link between two fixed points: a home and a small cell for example.

5G FWA could be either disruptive or complementary to cable providers. A telco or overbuilder with fiber capacity in a specific geographic area could use that fiber for backhaul and overlay a broadband radio area network (RAN) on top of the cable operator's service area to compete for broadband customers. This is clearly disruptive. Conversely, the cable operator could utilize 5G FWA to extend service to areas with no service or poor service without the expense of upgrading each premises and the surrounding infrastructure.

Powering an FWA network segment is similar to powering other HFC active network elements. Power per device, coax length between devices and distance from the power supply all factor into the powering equation. A simple example shown in Figure 10 illustrates the concept.



**Figure 10 – 5G Fixed Wireless Access Powering**

In this example a DAA node is shown servicing a neighborhood. Assume 4 small cell radios can service a total of 100 homes. To keep this example simple let's, make the following assumptions:

- The node and UPS power supply are co-located near a 4-way splitter
- Coax length from the splitter to each small cell is 1,500ft
- The UPS is configured to 90VAC output and is rated at 1350W
- The DAA optical node consumes 190W
- Each small cell requires 100W to power
- A gateway device interfaces each small cell radio to the coax for power conversion and DOCSIS backhaul. Each gateway requires 20W
- The coax is a common 0.625in diameter cable with resistance of 0.0011 ohms/ft
- The gateway / small cell minimum input voltage is 45V
- The operator's powering policy states that any broadband UPS can be loaded to maximum 85% of rated capacity.

With these parameters we calculate the power example as follows:

$$P(\text{Actives}) = P(\text{node}) + 4 \times (P(\text{small cell}) + P(\text{gateway}))$$

$$\text{Total Power (Actives)} = 190\text{W} + 4 \times (100\text{W} + 20\text{W}) = 670\text{W}$$

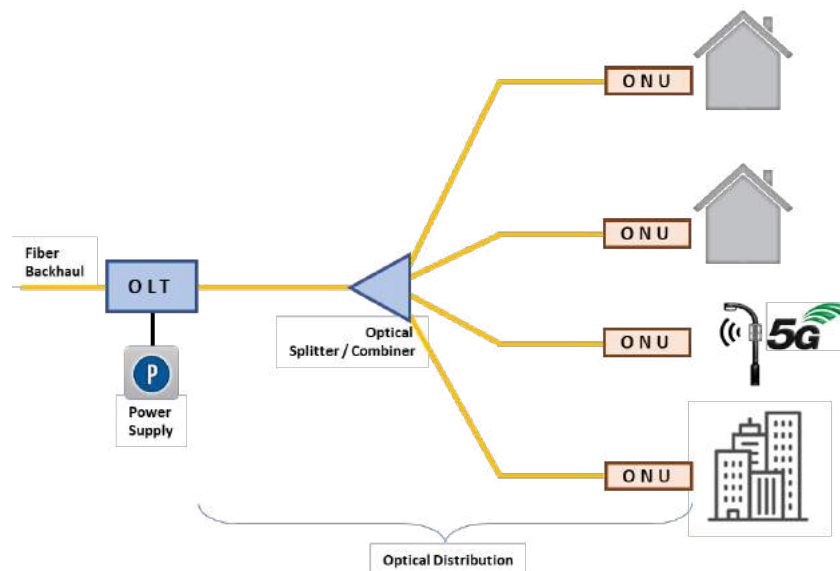
Adding coax line loss or  $I^2R$  loss adds an additional 13W. The results are summarized:

**Table 7 – 5G Fixed Access Wireless Powering Results**

| Configuration                                                      | PS I(out) | EOL Voltage | Actives Load | $I^2R$ Loss | PS Utilization |
|--------------------------------------------------------------------|-----------|-------------|--------------|-------------|----------------|
| DAA node, 4x 5G small cell radios each with HFC interface gateways | 7.6A      | 86.7V       | 670W         | 13.4W       | 51%            |

### 3.2.5. EPON

Ethernet passive optical network (EPON) has supported 10Gbps symmetrical data for years. Recently, the Institute of Electrical and Electronics Engineers (IEEE) has approved a 25G/50G-EPON standard. When viewed through the lens of the 10G initiative, cable broadband customers with EPON service are in good shape. Figure 11 represents an EPON network block diagram.



**Figure 11 – EPON Block Diagram**

Powering an EPON network segment is straightforward and familiar to most operators. EPON is an all-fiber network which consists of only passive optical components (splitters and combiners), except at the endpoints of each fiber, where there is an electrically-powered termination device – either an optical line terminal (OLT) or optical network unit (ONU). EPON is a point-to-multipoint topology in which downstream transmission from an OLT is received by all ONUs, but upstream transmission by an ONU is received only by the OLT. The OLT is often powered from a broadband UPS. That UPS may be dedicated to the OLT or may also power optical nodes and amplifiers in adjacent network segments. The ONU is powered from the premises (home or business) except for cases where the ONU services part of the access network infrastructure. An example is shown in our EPON block diagram where one ONU is

providing a backhaul connection to a 5G small cell radio. In such cases, the ONU and 5G small cell radio are typically powered from a nearby coax segment as described in a previous example.

## **4. Ensuring Reliable Power**

### **4.1. Power Reliability Considerations**

In 2019 a North American operator was experiencing problems with their digital nodes. Some of their R-PHY nodes were resetting. After a reset these nodes required several minutes to power-up and provision. The result was loss of service for several hundred customers within the service areas of the effected nodes. After extensive investigation by the operator and by multiple vendors the root cause for the resetting nodes was determined to be power related. Technicians working downline from the node were servicing network components and had inadvertently shorted the coax center conductor to ground causing a brief power disruption. The power anomaly was sufficient enough to cause the node to reset, dropping customer service for several minutes.

This example illustrates the critical nature of reliable power. Power reliability is especially vital with digital nodes where power disruptions can cause CPUs to reset. The node must reboot and then re-establish communications links independently to both the headend and to each user. The re-boot and re-provisioning cycle have been observed to take up to 15 minutes with some DAA devices. Many power related service disruptions can be mitigated with planning and preparation. A few ideas are discussed here.

#### **4.1.1. Intra-Node Power Hold-Up Time**

In the prior node reset example the cause of the power disruption was a momentary short of the coax cable carrying power to the node. A UPS system that is operating perfectly could not mitigate this type of power disruption. To the UPS system a coax line short would appear as a current spike on its output. If the short occurred close enough to the UPS, its ferro-resonant transformer would fold-back, dropping output voltage until the fault was cleared. If the line short was some distance from the UPS, the coax line resistance would mitigate the short and it would appear to the UPS output as a temporary rise in current.

The UPS has no way of protecting the node power input from this type of line fault. Avoiding power glitch related node resets requires keeping the node's internal logic power bus from dropping. Node vendors are experimenting with internal capacitors and batteries to this end. Capacitors should prove effective for short power disruptions of no more than a few 60Hz AC cycles (60-70ms). Extended hold-up times may require an internal battery or second power input to the node. Neither option is desirable. Internal batteries would require eventual replacement and a redundant second power input is complex and expensive.

#### **4.1.2. Utility Backup Time**

Increased power consumption from equipment such as future DAA nodes requires a review of UPS runtime capacity. With new architectures and equipment, what is the net effect on power consumption? Do the UPS systems still meet minimum runtime requirements? Operators must evaluate their networks by the criticality of each location. New DAA nodes may service business customers requiring longer utility backup than residential customers.

One utility power backup exception that California based operators must address is described in State Senate bill No. 431 introduced in June 2019. This bill requires some telecommunications equipment in



high fire threat areas to support 72 hours of communications during utility outages. If this legislation is applied to cable operators, extensive upgrades to their power backup systems will be required. For the access network, this would likely take the form of natural gas or propane generators at each broadband UPS installation. However, is it wise or even practical to have natural gas, propane or other combustible fuel source in zones declared high risk fire zones?

#### **4.1.3. Redundant Power Source**

Network equipment such as VHubs and DAA nodes often support large service areas or critical business customers. Do these devices have adequate power backup? In some installations operators have utilized a second (redundant) power pack within the node or VHub to insert power from a second UPS to create a fully redundant power backup for these critical devices. This redundant power may come from a second dedicated UPS system but would likely be diverted from a nearby coax segment powered from a different UPS than is powering the critical node or VHub.


### **4.2. Backup Power Requires Healthy Batteries**

The ability of broadband UPS systems to provide uninterrupted power is directly related to the condition of the UPS system batteries. Aspects of battery health and maintenance relating to outside plant (OSP) UPS operation are discussed here.

Multiple factors impact battery runtime and health. For clarity, runtime as used here is defined as the instantaneous stored energy available from a battery or bank of batteries irrespective of environment or history. Battery health is defined as the present maximum capacity of a battery or bank of batteries. Five primary factors affect battery runtime and battery state of health. These are: state of charge, ambient temperature, temperature history, battery age and charge history. Sections in 4.2 discusses charge history and its effect on both battery runtime and battery health or capacity. Note that this discussion applies to lead-acid types of batteries, which are commonly deployed in HFC networks due to their lower cost. As new battery technologies become available and cost effectively deployed, the conclusions below will need to be modified based on the characteristics of the newer battery technologies.

#### **4.2.1. Effects of Charge on Battery Capacity**

Overcharging and undercharging batteries has a significant effect on battery life. Using charging specifications from battery manufacturers, the following discussion illustrates how overcharging or undercharging batteries can cause premature battery failure.



**Figure 12 – Charge Effect on Battery Life**

In this diagram, the X-axis identifies specific voltages under and over the optimum charging voltage. The Y-axis shows the acceleration factor or multiplier on battery life. For example, if a battery were undercharged by 0.6V, its effective age would accelerate by a factor of 2x resulting in a reduction of useful life to 50% of that battery's optimum life. Likewise, if a battery were overcharged by 0.5V, its useful life would be only 30% of its optimum life.

Any charge related degradation effect occurs only while the overcharge or undercharge condition is applied. For example, undercharging a battery by 0.6V for a period of 4 months (perhaps the duration until the next preventative maintenance cycle) would result in a loss of 2 months to the overall life of the battery (using a 2x acceleration factor for 0.6V undercharge.) Once the charge problem is corrected, no further damage will occur, however, the capacity has nonetheless been permanently diminished.

#### **4.2.2. Battery Chemistry and Charge Mismatch**


Typical OSP power supply installations include one or multiple battery strings. Each battery string consists of 12V batteries connected in series. Three series batteries are combined to achieve a 36V string or four batteries are configured for a 48V string. The connection between the power supply and the battery string(s) is through a wire harness connected across the entire string (i.e., one connection is at the negative terminal of the first battery (ground), the second connection at the positive terminal of the third battery for 36V strings or the fourth battery for 48V strings). In this configuration, power from the batteries and charge to the batteries is routed through the wire harness for the entire string. A three battery, 36V battery string is shown in Figure 13.



**Figure 13 – Battery Charge Configuration**

Optimal float voltage is also listed for this 36V string as 13.62Vdc per battery. This value may vary among different battery manufacturers and technologies. The broadband UPS is configured to supply 40.86Vdc across the entire battery string. From the charger's perspective, each battery appears in the circuit as a fixed resistance. Per Ohm's law, the three batteries, acting as resistors in series, create a voltage divider and the 40.86Vdc is distributed equally across each of the batteries at 13.62Vdc per battery.

Batteries operate via electro-chemical reactions. Time, temperature and charge history can affect this chemistry, thus altering the battery's internal resistance. If each battery in a battery string has internal resistance values that change at different rates, the 40.86Vdc charge voltage will not be applied equally across each battery. The result of this unequal internal battery resistance is illustrated here.



**Figure 14 – Internal Battery Resistance**

As the internal resistance of each battery changes at different rates over time, the effective circuit where  $R1 \neq R2 \neq R3$  causes the charge voltage to be distributed unequally across the three batteries. The result is that some batteries will be undercharged while others are overcharged. In our example one battery is undercharged at 12.79Vdc while the other two batteries are overcharged at 14.20Vdc and 13.78Vdc. The total charge voltage of 40.86Vdc is correct, but the battery chemistry has caused an internal string charge variation that will shorten the life of all three batteries over time if not corrected.

OSP monitoring software can be configured to identify disparate voltages within individual battery strings for maintenance before extended battery damage occurs. Operators should consult battery manufacturers for specific voltage threshold parameters to trigger alerts and initiate corrective action. Multiple options exist to mitigate the effects of time on internal battery resistance. Some modern UPS systems are equipped with charge balancing technology that will automatically re-direct charge within a battery string to offset the effect of changes to internal battery resistance. This charge management technology is available from various manufacturers in a variety of configurations. Operators should be aware of the effects of charge imbalance on battery life and determine the best course of action for their situation.

#### **4.2.3. Preventative Maintenance is Essential**

In 2016, a major North American operator conducted an investigation to determine the root cause of their growing number of OSP broadband UPS alarms. Standby test fail alarms were of particular interest due to the critical nature of this alarm. Across three (3) cities, this operator identified that 22% of their broadband UPS systems had failed standby tests, due in part, to battery cable corrosion. The following shows representative battery corrosion at one of these installations.



**Figure 15 – Battery Corrosion**

The center battery shows excessive corrosion around both the positive and negative battery posts. This type and level of corrosion can occur with some styles of OSP batteries that have not been properly maintained. At the conclusion of their investigation, the operator identified deficiencies in their local OSP preventative maintenance (PM) practices that led to these potentially service-impacting results.

Corrosion will damage battery power cables and battery voltage monitoring sense wires. As battery power cables degrade, electrical resistance increases and the ability of the batteries to provide sufficient current diminishes. Eventually, as this operator experienced, the corrosion will increase cable resistance sufficiently to cause a broadband UPS standby test to fail. Prior to a standby test failure, backup capacity had diminished and the expected runtimes during actual power outages were lower. Had an actual utility

outage occurred, standby power would be compromised and customer loads dropped, potentially before the standby test indicated any problem.

Could this liability have been avoided through more diligent preventative maintenance practices? The answer, of course, is yes. “How frequently should each UPS be visited?” is an often debated question. Responses vary between operators and even between systems within the same operator. These answers range from six-months to two-years when technicians are queried. Conducting PM visits and finding nothing to correct is wasting valuable service resources. Waiting too long between PM visits could result in service impacting situations going unchecked. One often hears reports of PM visits to multiple installations within a geographic area with some locations checking out OK while other sites require extensive maintenance. Clearly, there is no one right answer to the question of PM frequency. Can anything be done to reduce unneeded PM visits while focusing limited resource on locations needing physical intervention?

The answer to this question is two-fold:

- First, there are some issues that require on-site inspection to identify and correct. Examples include pest infestations and water intrusion due to physical enclosure damage. Because this category of problems exists, scheduled preventative maintenance visits are required.
- Second, it may be possible to identify a category of future service-affecting problems through analysis of data available from status monitoring systems. This would enable operators to prioritize site visits around locations at high risk of causing future service disruptions. Low risk sites could be visited less frequently and only to inspect and correct issues that are undetectable any other way.

## 5. Conclusion

This white paper discussed access network powering considerations related to network upgrades intended to support new services that will reach and eventually surpass performance goals of the 10G initiative.

A traditional HFC network powering approach was used as a baseline for other powering examples. Upgrading to a DAA network topology was reviewed and a near future DAA node defined and used in several powering scenarios. Between the DAA node and the premises, powering for multiple network architectures was discussed including last-mile access via a deep fiber N+0 approach and a 5G fixed wireless access approach.

Several powering themes became apparent as different access network scenarios were reviewed.

1. Experience with powering of traditional HFC networks is foundational. An understanding of both active loads and  $I^2R$  losses is a good starting point for analyzing power demand of network upgrades. Coax line losses play a major factor in powering decisions, especially as new active devices are placed some distance from the UPS.
2. Digital line gear including DAA nodes require more power than their analog counterparts and they are more sensitive to power disruptions. A DAA node (or VHub) reset has the potential to drop many customers for several minutes as processors reset and communications are re-provisioned. Policies regarding power quality, redundancy and backup time should be established before any significant DAA system upgrades.
3. Since access network equipment presents a constant power load to the network,  $I^2C$  line loss and EOL voltage are interrelated. Higher line loss (due to coax length as well as other factors) results in more current draw through the coax which results in lower EOL voltage.

This paper also discussed ways to ensure power reliability with an emphasis on batteries. Battery charge management is essential to battery life and ultimately to maintaining network uptime during power utility disruptions. Internal battery chemistry changes over time will result in the battery internal resistance changing at different rates, requiring active monitoring and management to ensure that batteries are charged correctly over their entire service life. Finally, on-site maintenance is needed to oversee aspects of the UPS system that cannot be managed remotely, such as damage and corrosion from weather and fauna.

This paper reviewed only a few representative network powering scenarios. Powering designs for live networks must be engineered alongside the network architecture being powered to ensure that power quality and quantity is sufficient for the near future network requirements.

## Abbreviations

|          |                                                                    |
|----------|--------------------------------------------------------------------|
| 10G      | 10 gigabit                                                         |
| 5G       | fifth generation technology standards for cellular networks        |
| AP       | access point                                                       |
| ASIC     | application specific integrated circuit                            |
| bps      | bits per second                                                    |
| CBRS     | citizens broadband radio service                                   |
| CCAP     | converged cable access platform                                    |
| CPU      | central processing unit                                            |
| DAA      | distributed access architecture                                    |
| DCA      | distributed CCAP architecture                                      |
| DOCSIS   | data over cable service interface specification                    |
| EPON     | Ethernet passive optical network                                   |
| ESD      | extended spectrum DOCSIS                                           |
| FDD      | full duplex DOCSIS                                                 |
| FEC      | forward error correction                                           |
| FPGA     | field programmable gate array                                      |
| FWA      | fixed wireless access                                              |
| GAP      | generic access platform                                            |
| HD       | high definition                                                    |
| HFC      | hybrid fiber-coax                                                  |
| Hz       | hertz                                                              |
| IoT      | internet of things                                                 |
| ISBE     | International Society of Broadband Experts                         |
| MAC      | media access control                                               |
| N+0      | node plus zero amplifiers                                          |
| ODC      | optical distribution center                                        |
| OLT      | optical line terminal                                              |
| ONT      | optical network terminal                                           |
| OSP      | outside plant                                                      |
| OTT      | over the top                                                       |
| PIM      | protective interface module                                        |
| PNM      | proactive network maintenance                                      |
| PON      | passive optical network                                            |
| RAN      | radio area network                                                 |
| RF       | radio frequency                                                    |
| R-MACPHY | remote mac and physical layers                                     |
| R-PHY    | remote physical layer                                              |
| SCTE     | Society of Cable Telecommunications Engineers                      |
| vCCAP    | virtual CCAP                                                       |
| vCMTS    | virtual cable modem termination system                             |
| Wi-Fi 6  | sixth generation wireless intent standards, also known as 802.11ax |

# Bibliography & References

- [1] Cable Labs, "10G: The Next Great Leap in Broadband," 2019.
- [2] CableLabs, "DAA (Distributed Access Architecture)," [Online].
- [3] CableLabs, "Distributed CCAP Architecture Overview Technical Report, CM-TR-DCA-V01-150908".
- [4] CableLabs, "Data-Over-Cable Service Interface Specifications, DOCSIS® 4.0, Physical Layer Specification, CM-SP-PHYv4.0-I02-200429".
- [5] J. Chapman, "Blueprint for 3 GHz, 25 Gbps DOCSIS® Getting 25 Gbps PON-Like Performance Out of HFC," *SCTE, Cable-Tech Expo*, 2019.
- [6] CableLabs, "P2P Coherent Optics Architecture Specification P2PCO-SP-ARCH-I02-190311".
- [7] SCTE, "Generic Access Platform Requirements, IPS WG1 (Draft)".
- [8] CableLabs, "Adrenaline Distributed Compute Platform," in *Cable Next-Gen Europe Digital Symposium*, 2020.



## **From CSP to DSP**

### **Is the COVID-19 crisis a partner or another steppingstone?**

A Technical Paper prepared for SCTE•ISBE by

**Javier Ger**

Cloud Infrastructure Strategy Manager  
Telecom Argentina  
Aguero 2392 - C1425EHZ, CABA - Argentina  
+541155304531  
jger@teco.com.ar

**Claudio Saes**

Regional Partner  
Bell Labs Consulting  
3100 Olympus Blv, Coppell - Texas 75019  
+1 (214) 208-4970  
claudio.saes@bell-labs.com

# Table of Contents

| Title                                             | Page Number |
|---------------------------------------------------|-------------|
| 1. Introduction.....                              | 4           |
| 2. General Pre-COVID Context.....                 | 4           |
| 3. CSP's Pre-COVID Context.....                   | 5           |
| 3.1. CSP's Challenges.....                        | 5           |
| 3.2. CSPs need to do things differently.....      | 6           |
| 3.3. What does transformation into DSP mean?..... | 7           |
| 4. COVID-19 pandemic impacts.....                 | 12          |
| 4.1. Global Economy.....                          | 12          |
| 4.2. CSPs Traffic.....                            | 13          |
| 4.3. Other CSPs Aspects.....                      | 15          |
| 4.4. Some key questions remain.....               | 17          |
| 4.5. Pandemic Opportunities.....                  | 17          |
| 5. Post COVID-19 world.....                       | 18          |
| 5.1. CSPs Additional Opportunities.....           | 18          |
| 5.2. Three probable hypothesis.....               | 21          |
| 5.3. CSPs Decisions.....                          | 22          |
| 6. Conclusion.....                                | 22          |
| Abbreviations.....                                | 26          |
| Bibliography & References.....                    | 26          |

## List of Figures

| Title                                                                                                             | Page Number |
|-------------------------------------------------------------------------------------------------------------------|-------------|
| Figure 1 – CSPs vs hyperscale current scenario outlook.....                                                       | 5           |
| Figure 2 - Nearly US\$9B untapped industry 4.0 value to be addressed by CSPs (McKinsey Global Institute).....     | 7           |
| Figure 3 – CSPs transformation options.....                                                                       | 7           |
| Figure 4 – Six tier 1 CSPs' EV/EBITDA, ROC and NPS as of December 2019.....                                       | 8           |
| Figure 5 - The four digital transformation pillars.....                                                           | 9           |
| Figure 6 – The “S Transformation” approach.....                                                                   | 10          |
| Figure 7 – Kotter's vision about Digital Transformation Pitfalls.....                                             | 11          |
| Figure 8 – Global Economy Projections.....                                                                        | 13          |
| Figure 9 – Global Pandemic Traffic Scenario.....                                                                  | 14          |
| Figure 10- Higher traffic volumes and peak hour period extension.....                                             | 14          |
| Figure 11 – Public traffic impact from Telecom Argentina, April 2020.....                                         | 15          |
| Figure 12 - Historical CAPEX/Revenue ratio vs GDP.....                                                            | 16          |
| Figure 13 – Some B2B/B2B2C use cases with strict SLA – Bell Labs Consulting.....                                  | 18          |
| Figure 14 – Edge Battle – Footprint granularity can be a CSPs valuable asset coupled with the right strategy..... | 19          |
| Figure 15 - 5G announced activities by industry OMDIA July 2020.....                                              | 20          |
| Figure 16 - Digital Transformation's hypothesis after COVID-19.....                                               | 21          |

## List of Tables

| <b>Title</b>                                                              | <b>Page Number</b> |
|---------------------------------------------------------------------------|--------------------|
| Table 1 - 5G public announcements for enterprise - OMDIA, July 2020 ..... | 20                 |

## 1. Introduction

The COVID-19 health crisis is one of those unique transformative moments in time, resembling the effects of wars, economic, political, or governmental crisis. After a few months, human beings are trying to adjust to the changes triggered by social distancing effects.

Customer service providers (CSPs) are in the center of this cataclysm, providing the connection between people, entertainment, retail, hospitals, and governments, accelerating the need for more significant usage of digital services. The major part of the CSPs is embracing the journey to becoming a digital service provider (DSP). Being a DSP will require providers to attend to dynamic customers' demands. It will need to be part of a broader ecosystem of participants to offer not only connectivity but a full suite of digital products and agile services to end customers and partners with a complex business model—and finally increasing their value share in the market.

However, as COVID-19 lockdowns extend, the impacts in the macro-economies are profoundly affecting the population and the CSP's cash flows, forcing them to limit or even stop their investments for a few months. This change management process is shifting the CSP's focus to a myopic investment mode and adding to it telecom industry supply chain disruption shortage, this scenario is the recipe for stagnation.

Against all the odds, some CSPs are finding the creativity and the innovation to meet their customer's needs, but still far from the disruption model proposed by the digital transformation initiatives.

In this paper, the authors will examine the impacts and trade-offs of COVID-19 into societies, industries, economy, and the information and communications technology (ICT) market, reflecting on some post-COVID initiatives CSPs should embrace to expedite the informal labor economy, meet new digital consumer habits, attend public safety requisites and, definitively, reestablish its course to a sustainable digital service provider path.

## 2. General Pre-COVID Context

New technologies have enabled brilliant innovations and the new kid on the block is called the Fourth Industrial Revolution. According to the World Economic Forum<sup>1</sup>, this revolution is defined by a fundamental change in the way we live, work, and relate to one another by merging the physical, digital and biological worlds. In its scale, scope, and complexity, the transformation will be unlike anything humankind has experienced before, in ways which we never saw before. One of the fundamental differences of the current industrial revolution from the previous ones is that previously the changes occurred outside of the human body; instead, in this case, they will involve it and even happen within it.

Today, industries around the world like factories, mining, logistics, and others have embraced their digital transformation. With the introduction of new technologies, widespread connectivity, and cloud workloads, these industries will transform themselves – increasing their process efficiencies, workforce safety, and growing their productivity – generating greater value to the market and back to the communities.

Beyond the current and known initiatives, the next ten years promise us a future where humans will be augmented by artificial intelligence (AI) and digital, physical and biological systems that will work in a tightly coupled mode, possibly changing us, as homo sapiens.

---

<sup>1</sup> <https://www.weforum.org/focus/fourth-industrial-revolution>

To name a few examples of the expected human revolution, we may see:


- Disappearing boundaries for human beings between what is natural and what is artificial. Not only as part of the human body itself but including concepts as “mirror world” or metaverse<sup>2</sup>.
- Changes between our relationships with life, the planet, work, and consumption. The industries are aiming to cover every human aspect need and generating a new history of how we live.
- Thus, it covers aspects related to health care, equality, materials, energy, political and economic systems, having a context to produce a sustainable life for the planet and human beings.
- This comprehensive set of aspects proposes working with new and natural materials to decouple the economic growth from the constraints of the existing raw materials.

### 3. CSP's Pre-COVID Context

#### 3.1. CSP's Challenges

CSPs have been facing a hard competitive scenario due to emerging digital disruptors, who are offering higher-value services demanded by the consumers and leveraged on the CSP's infrastructure. As most CSPs compete for similar business to consumer (B2C) & business to business (B2B) services with the lowest competitive advantage, revenues are commonly disrupted by new players. This disruption is commonly based on a combination of platforms, experience, and costs, reducing dramatically the friction for customers to access the products and services they are looking for.

This scenario is shown in Figure 1 and represents something well-known in the industry. The CSP's revenues are becoming flattened and at the same time, their total cost of ownership (TCO) has increased drastically by significant infrastructure deployment to manage the traffic growth generated by hyperscalers and new consumers' demands.



**Figure 1 – CSPs vs hyperscale current scenario outlook**

<sup>2</sup> <https://es.wikipedia.org/wiki/Neuralink>

<https://www.forbes.com/sites/cathyhackl/2020/07/05/the-metaverse-is-coming--its-a-very-big-deal/#3cc83452440f>

[https://en.wikipedia.org/wiki/Mirror\\_world](https://en.wikipedia.org/wiki/Mirror_world)

### 3.2. CSPs need to do things differently

In the past five years, the telecom industry has been one of the worst-performing sectors for investors as global bets have failed to pay off. That has left companies with huge debts, and now CSPs are under intense pressure to invest in new 5G and full-fiber networks.

On the other hand, regulation has also played a significant role in challenging the telecom business model's execution. Meanwhile, OTTs are scoring higher revenue streams as they have better control of data propelling its business forward. Still, the difference in OTT regulation compared with telecoms has contributed to some disparity of outcomes.

Disruption in technology takes place silently, often too late to be noticed or worst reacted. History has proven that businesses that failed to perceive disruption and re-invent themselves had lost a tremendous amount of value or disappeared- and we all know who they are.

The foundational question is why CSPs fail to perceive the disruption surrounding them?

Kahneman, Lovallo, and Sibony<sup>3</sup> had provided us useful insights on complex corporate decision-making bias; For instance, preconceived notions are difficult to be contradicted even when confronted with strong evidence, leading to an effect named confirmation bias.

Anchoring decisions do leaders weigh one piece of information too slowly in making decisions, and finally, loss aversion makes them too cautious.

Most of today's CSPs' cash flows are defensive and, for the most part, structured by a supplier-buyer mindset and the blame game when things go wrong.

Suppose the traditional telecom operators don't try different strategies. In that case, they will lose value on the market and potentially be acquired, or struggle for the next few years with average growth, similar to the historical examples.

Based on the presented outlook, CSPs need structural transformations on their businesses to be able to compete for higher-value products and services, allowing them to reduce the gap between the revenues and traffic, cover the new and extremely dynamic consumers' demands in this fluctuating world and be able to make their business more sustainable.

**GAP + New Demands + Competition = New Products + Improve TTM<sup>4</sup> + Optimize TCO**

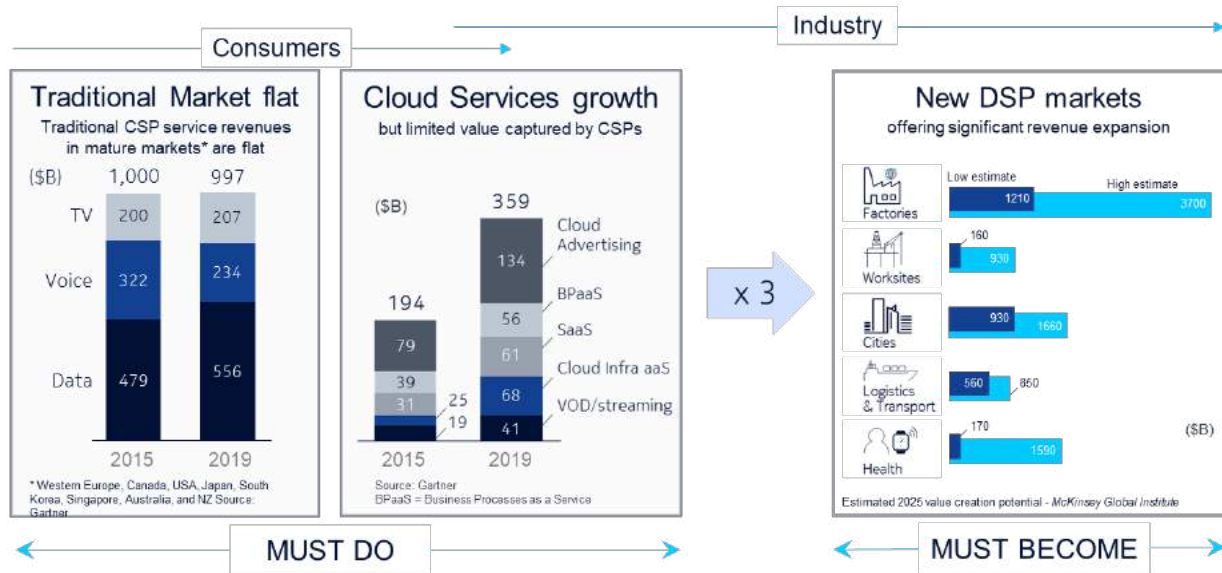
#### *Transformation*

If well inserted and with the right strategy, CSPs have the potential to disrupt markets, inserting themselves into new value chains & capturing new revenues from untapped industries going through Industry 4.0 development, which promises to bring more value in the short-mid term. According to the McKinsey Institute<sup>5</sup>, there's near US\$9T untapped value to be addressed in this sector by while the largest segment of traditional CSPs market ranges about US\$1T - Figure 2.

<sup>3</sup> Kahneman, D., Lovallo, D. and Sibony, O., 2011. The Big Idea: Before You Make That Big Decision.... [online] Harvard Business Review

<sup>4</sup> Time to Market


<sup>5</sup> <https://www.mckinsey.com/business-functions/mckinsey-digital/our-insights/the-trillion-dollar-opportunity-for-the-industrial-sector>



**Figure 2 - Nearly US\$9B untapped industry 4.0 value to be addressed by CSPs (McKinsey Global Institute)**

### 3.3. What does transformation into DSP mean?

Eying the benefits described in the previous sections, many CSPs started journeys to become themselves more agile and digital companies, namely digital service providers or DSPs. According to Martin Creaner<sup>6</sup>, DSPs are those capable of responding to the changing demands of their clients, fostering and being part of larger ecosystems based on much more complex business models, offering not only connectivity but a wide spectrum of digital products and agile services and therefore increasing the value for their end customers, partners and shareholders (see Figure 3)



**Figure 3 – CSPs transformation options**

The authors researched six global CSPs located in six unique countries (see Figure 4). This picture emphasizes the difference in scale between current CSPs and hyperscalers and, for this reason, the need

<sup>6</sup> Creaner, M., 2020. *Transforming The Telco: Developing The Instinct For Survival In The Digital Economy*. CenterNODE.org.

<https://www.youtube.com/watch?v=xl6gIEb8OAM> (Martin Creaner - Former CEO & President - TM Forum)


for this CSPs transformation to achieve greater value, allowing them to increase their margins in this new scenario.

In other words, those CSPs deciding to close the gap must offer higher value-added services and products to consumers. This will be only possible by implementing structural transformations, which will require strong leadership to manage the risks and failures.

In Figure 4, the listed CSPs are going or went through serious digitalization efforts, supported by automation, big data analytics, machine learning (ML), AI, network function virtualization (NFV), software-defined networks (SDN), and cloud strategies ranging from two to eight years.

These initiatives aim to improve customer experience, provide operational expenditure (OPEX) savings, and allow a shorter time to market (TTM). However, after careful investigation, these CSPs' digital transformation initiatives are failing to deliver on their value and efficiency results according to public market information.

Their valuation over earnings before interest, taxes, depreciation, and amortization (EBITDA) varies quite considerably, as well as their perceived Net Promoter Score (NPS), providing signs that it takes a good strategy and organization alignment to generate value back to the investors and clients.



Sources: December, 2019 - S&P, prof Damodaran Stern School of Business, Net Promoter Score Guru

**Figure 4 --Six tier 1 CSPs' EV/EBITDA, ROC and NPS as of December 2019**

An effective approach to transformation requires establishing the business vision, assessing the present, and guiding the transformation to the target state. Aligned with its business strategy and a clear transformation roadmap, a CSP will need to re-arrange its organization's culture, skills, process, and technology to successfully embrace its digital transformation initiatives.



A common misconception is that if a company deploys an extra set of technologies, then digital transformation is done. However, Becky Frankiewicz and Tomas Chamorro-Premuzic<sup>7</sup> explain that the reality is that digital transformation isn't just about technology, but it's about organizational adaptability. To keep pace with the changes driven by digital transformation, organizations must learn to be an agile, adaptable, and organizational culture; and talent skills are crucial to the success of any digital initiative, including unlearning old behaviors.

For genuine change to happen, company-wide support is critical. The CSP leadership and executive teams must provide a clear vision and strategy for how the change will be realized, prioritize it properly, define concrete goals and accountability for the firm and measure its achievements. Thus, the entire organization must point in the right direction. The employees must understand their role in the transformation and how they are driving change. New technologies and change management will require new skills, culture, and processes, too.

The truth is that CSPs will need to have aligned the four pillars of the business strategy shown in Figure 5, the key structure over which workforce skills and culture, process, and technology will be based, to create a framework to support their employees to deal with the new software-centric technologies.



**Figure 5 - The four digital transformation pillars**

To promote cultural cohesion, soft-skills training will need to be included, intellectual curiosity must be permanent and fostered for continuous learning. Other important characteristics are teamwork, adaptability, and flexibility to unknown conditions, change management practices, effective communication, negotiation, and conflict resolution.

Re-skilling can cost a reasonable amount of financial investments and will need to be surgically inserted in the operator's context and budget to provide the right outcomes.

A business strategy, described above, is the fundamental pillar of the digital transformation. One typical proposal is to move forward with operational efficiency optimization, improving the current businesses as

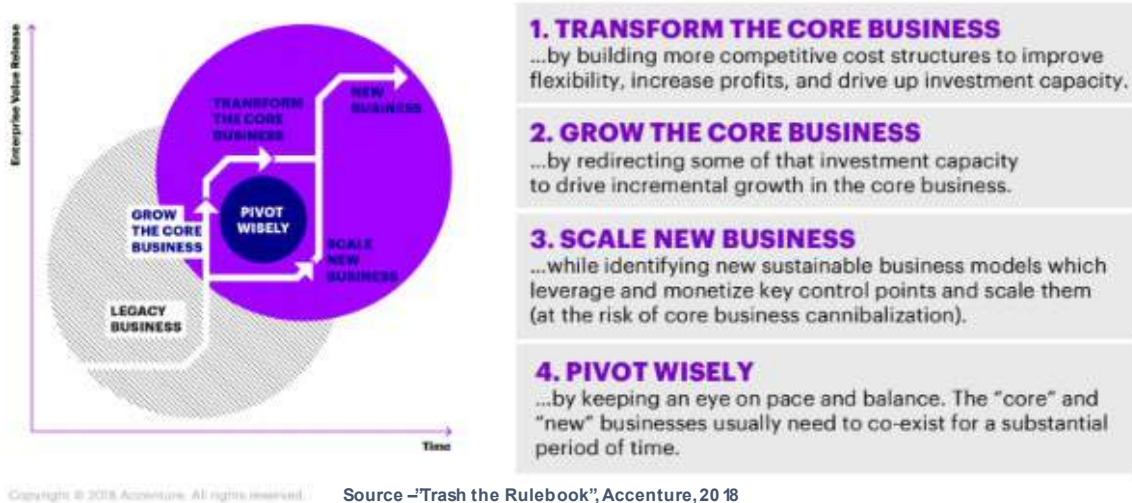
---

<sup>7</sup>“Digital Transformation Is About Talent, Not Technology”, Becky Frankiewicz and Tomas Chamorro-Premuzic, Harvard Business Review, May 2020. Available online at: <https://hbr.org/2020/05/digital-transformation-is-about-talent-not-technology>.

much as possible. In contrast, new services, products, and business models should be investigated to leverage competitive advantage. This approach is defined as the "S Transformation"<sup>8</sup>- Figure 6.

*The transformation process is not a project; it does not have an end; conversely, it is something that once it starts is the new mode in which the company will work. It becomes a new Business as Usual.*

## SCALING NEW GROWTH IS A DELIBERATE AND PERPETUAL CHANGE JOURNEY, NOT A SINGLE EVENT



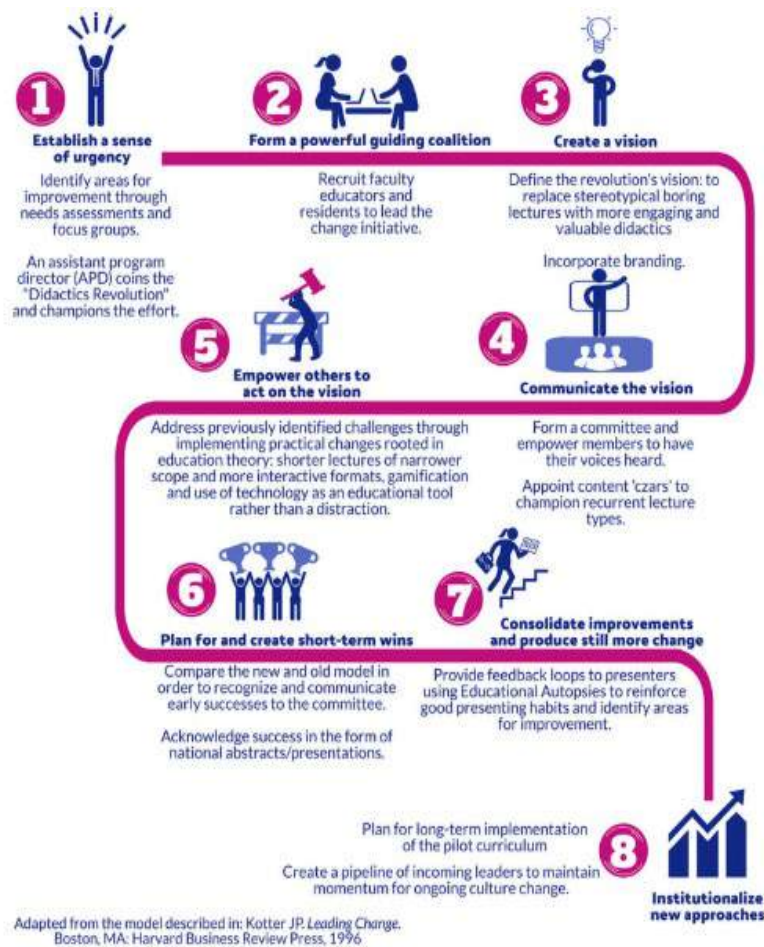
**Figure 6 – The "S Transformation" approach**

As CSPs are looking for ways to move towards its digital transformation journeys, there are notable pitfalls they should avoid.

Much has been studied and written on this topic and these extensive projects seem to have an interesting commonality on its pitfalls. According to the authors, one of the most comprehensive analysis of this topic is covered by John Kotter<sup>9</sup>. Figure 7 describes a graphical and rather pedagogical version of his framework.

<sup>8</sup> Accenture, 2018. Trash The Rule Book.

<sup>9</sup> "Leading Change - Why Transformation Efforts Fail", John P. Kotter, Harvard Business Review, January 1996



**Figure 7 – Kotter’s vision about Digital Transformation Pitfalls**

According to another study, Mike Sutcliffe, Raghav Narsalay, and AaroHi Sen<sup>10</sup>, there are two major digital transformation challenges observed:

- 1) A silent disagreement between top-level managers on the project goals, leaving their direct reports confused on what to prioritize, and how to measure progress.
- 2) A gap between the digital capabilities that were used to support a pilot, and the capabilities required to support the scaling of it. If this problem isn’t addressed, the firms could face lengthy delays to scale-up production or leadership team’s hurrying things to meet the promised changes.

Both could be remediated, but they require communication and negotiation skills. The first challenge needs to be discussed between the stakeholder parties. The opportunities and the benefits need to be articulated, but also the problems it solves and how the firm will re-orient the organization before investing in the desired solution.

<sup>10</sup> “The Two Big Reasons That Digital Transformations Fail”, Mike Sutcliffe, Raghav Narsalay and AaroHi Sen, Harvard Business Review, October 2019. Available online at: <https://hbr.org/2019/10/the-two-big-reasons-that-digital-transformations-fail>.

The second challenge is more complex and will require the stakeholders to look outside the organization to close the divide or raise the capabilities from inside, starting with the pilot and growing from it.

Finally, Jacques Bughin and Nicholas Van Zeebroeck<sup>11</sup> describe six strategies to face this transformation. Within their studies, they concluded three offensive and three defensive options, being the former more successful to achieve the desired transformation.

- Offensive: Platform play, New marginal supply, Digitally-enabled products, and services
- Defensive: Rebundling and customizing, Digital distribution channels, Cost efficiency

Given that the pandemic crisis forced many operators to fix some operational inefficiencies, a few CSPs claim it helped them to prioritize their digital transformation efforts. We will look more carefully at this topic in the following sections.

*“Wisdom consists of knowing how to distinguish the nature of the trouble,  
and in choosing the lesser evil”  
The Prince, Niccolo Machiavelli*

## **4. COVID-19 pandemic impacts**

### **4.1. Global Economy**


In March 2020 many governments started taking measures to protect and mitigate the pandemic impact on the health of their populations. These decisions have heavily affected the global macroeconomy, as the International Monetary Fund (IMF) and World Bank have shown in their reports and forecasts - Figure 8.

At the writing of this paper, it is impossible to estimate the economic impact of the health crisis, as it depends on its duration. Although some forecasts provide a more optimistic scenario, according to the World Bank, this is the most significant recession since World War II.

The critical question is how the macroeconomy, political tension, and exchange rate variations will affect the CSP's investments and, in particular, how it will affect its digital transformation initiatives to address the 4<sup>th</sup> industrial revolution requirements?

---

<sup>11</sup> “6 Digital Strategies, and Why Some Work Better than Others”, Jacques Bughin and Nicholas Van Zeebroeck, Harvard Business Review, July 2017



**Figure 8 – Global Economy Projections<sup>12</sup>**

There were severe decisions taken by governments to control the spread of the disease, which are affecting several industries around the globe. A few examples are<sup>13</sup>:

- During the lockdown -stay-at-home, shelter-in-place, and quarantines-, nearly 3 billion population were living in countries whose borders were shutdown to nonresidents, and 93% of countries had imposed immigration limits.
- As of April 10, governments across the globe had announced stimulus packages amounting to \$10.6 trillion—this is the equivalent of eight Marshall Plans.
- The global macro-economy will affect enterprises, including CSPs, and finally, the governments will need to fund these stimulus packages to move the economy.

## 4.2. CSPs Traffic

Government measures described in the previous section have produced some changes in population habits. Work, study, leisure, and consumption from home are the new norm, all made possible by CSPs infrastructure around the world.

Change in traffic profiles has been significant, and behavior varies between regions and countries depending on several factors such as isolation intensity, lockdown periods, e-learning availability, and e-commerce maturity.

One remarkable aspect during the lockdown is that traffic volumes varied, but it all had a severe impact during peak hours affecting the current networks and the capacity planning cycles.

<sup>12</sup> <https://www.imf.org/en/Publications/WEO/Issues/2020/06/24/WEOUpdateJune2020>  
<https://www.worldbank.org/en/publication/global-economic-prospects>

<sup>13</sup> “The future is not what it used to be: Thoughts on the shape of the next normal”, Kevin Sneader and Shubham Singhal, McKinsey & Company, April 2020. The exact value of some figures from this article can be discussed, but undoubtedly, they are really large



The periods of peak hours had extended considerably, with upstream traffic struggling access network resources.

This can be seen in Figure 9, Figure 11, and Figure 11.

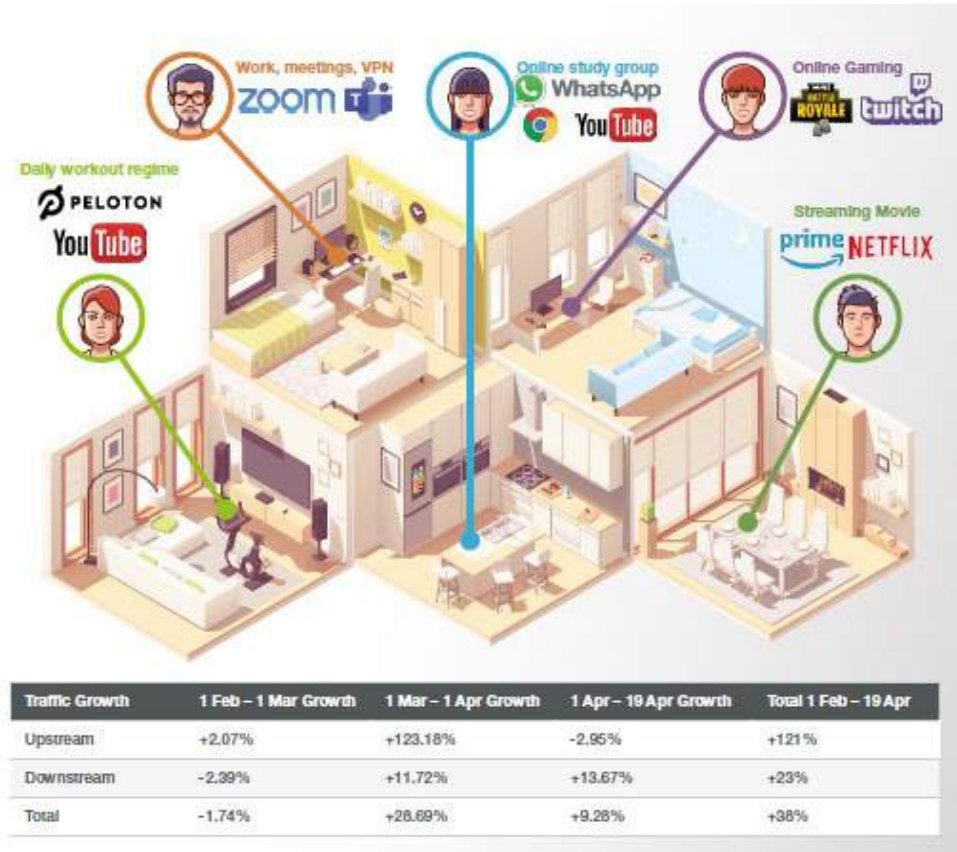


Figure 9 – Global Pandemic Traffic Scenario


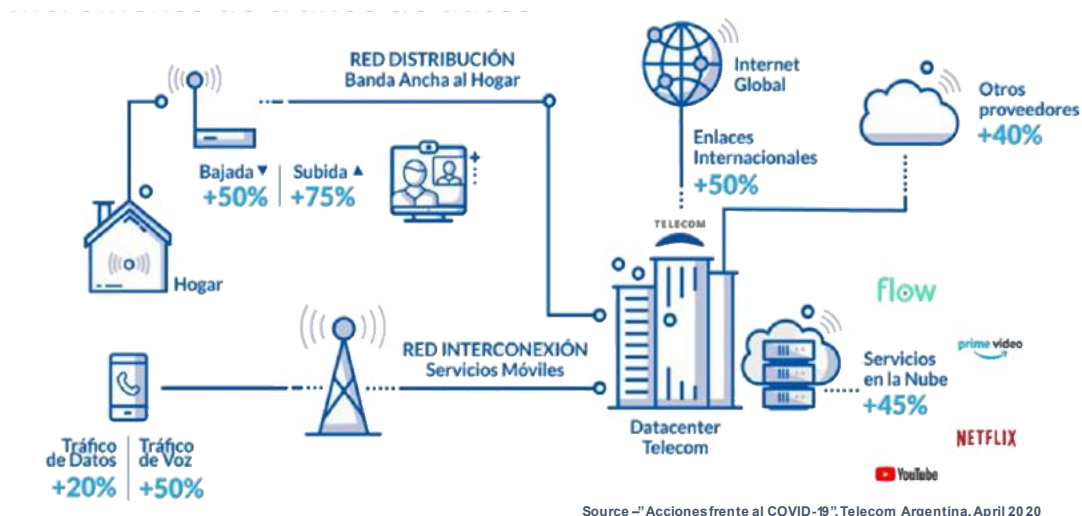


Figure 10- Higher traffic volumes and peak hour period extension



**Figure 11 – Public traffic impact from Telecom Argentina, April 2020**

### 4.3. Other CSPs Aspects

The situation was uncommon, but COVID-19 forced the CSPs to speed-up their operational initiatives to address the unexpected conditions.


During the past 5 months of social distancing and lockdown we observed:

1. CSP's networks have responded reasonably well to the change in traffic demands, including 40% to 50% traffic spikes.
2. Remote working and e-learning have amplified the gaps in the upstream capacity and also in the rural broadband connectivity.
3. The higher adoption of remote work accelerated the demand for the enterprise digitalization, including medium and large enterprises embracing more cloud services, and some CSPs providing higher capacity and resilient residential broadband access.
4. CSPs across the globe are going beyond the call of duty to help communities and consumers to manage COVID-19, including free access to e-learning resources, movies, games, and ultra-reliable services to first responder's authorities.
5. As consumers and SMEs' incomes were severely affected by economic stagnation and increasing unemployment rates, CSPs are facing a severe impact on their revenues and investment plans. Some consumers can't afford their billings, and some governments have temporarily suspended the option to disconnect the service.
6. IT Systems -e.g., business support systems (BSS) & operations support systems (OSS) - have to be re-oriented to these new conditions. In some countries, operators had to revise their prepaid recharge systems and policies quickly to provide a more seamless customer experience.
7. To keep up with the higher demands for traffic growth, e-learning, remote working, virtual private network (VPN) and security, CSPs had to revise their budget plans and invest capital expenditure (CAPEX) and OPEX to minimize operational issues.
8. Global supply chain disruptions have impacted the CSPs' ability to scale-up deployments.
9. On the fixed access networks, CSPs re-arranged their field operations to service their consumer's homes when outages related to their customer premises equipment (CPE) or residential gateways (RGWs) happened.

When CSPs look at the bright side of the COVID-19 pandemic, they list several opportunities to speed up a few digital transformation initiatives, such as those listed above.

As was previously described in this paper, also with the social distancing, the global macroeconomy is heavily affected and thus CSPs revenues and investments are at risk. When we look at the global CSP's revenue breakdown, approximately 50% of it comes from mobility products and services. As the population is in lockdown, they rely more on fixed internet access, affecting largely the CSP's mobile services and handset revenues.

During the 2nd and 3rd quarter of 2020, operators revised their CAPEX plans, and given the general economic and network effect from the pandemic, including revenue impact, disruption of the supply chain and the need to expand network infrastructure, their investments not related to short-term aspects are likely to fall compared to previous years.



**Figure 12 - Historical CAPEX/Revenue ratio vs GDP**

Although traffic has surged and profile has changed due to COVID, operators will focus on improving and maximizing existing network and technology, so other investments will be pushed to late 2021-2022 whenever possible

In many emerging countries we observe a delayed investment in 5G networks, while developed countries slowed down their 5G network deployment plans in the first half of 2020. However, similar to the 4G deployments, 5G is expected to ramp up from 2021 onwards. Additionally, given that most of the wireless CAPEX comes from radio access network (RAN) investments, and the population will tend to be indoors, investment in RAN will drop.



In the particular case of Latin America (LATAM) -see Figure 12- major CSPs will see a revenue decline in 2020 between -3% to -14%<sup>14</sup>, directly impacting CAPEX growth, but although gross domestic product (GDP) impacts revenue, capital intensity has been reasonably constant, with peaks mainly in wireless due to technology refresh.

#### **4.4. Some key questions remain**

CSPs around the world have been relatively successful in dealing with the challenges that the current pandemic is presenting. Many of them have been displayed in the previous section.

But some key questions remain.

- Was the CSPs' success in managing the COVID-19 crisis achieved with the same principles required to transform them into a DSP or did accelerated efficiency only follow the existing CSPs' operating model?
- Can these results be claimed as part of a digital transformation?
- Will the macro-economic outlook delay non-critical investments like 5G and transformations to DSPs?

#### **4.5. Pandemic Opportunities**

The previous questions are even more crucial considering the opportunities that can emerge from this crisis. Some of them can imply permanent changes and, therefore, opportunities to be taken in the short-term.

In areas like e-commerce, fintech, telemedicine, automation, remote working, and e-learning, the COVID-19 pandemic could be a decisive turning point in the short and mid-term, laying a foundation for digital services adoption for the future<sup>15</sup>. For example:

- In China, individuals aged 36 and over and residents of smaller, less prosperous cities have begun to shop online.
- In Europe, 13% of consumers said they were planning to move to online purchasing options.
- In Italy, e-commerce transactions have risen 81% since the end of February.
- The business opportunities for banks and fintech companies are exponential with shifts to digital payments and remote cashless transactions.
- In the US, one of the largest stand-alone telemedicine service providers reported a 50 percent increase in remote consultation during the pandemic.
- France, Korea, and Brazil changed its regulations to ease access to telemedicine.
- Automation efforts are increasing in every industry. According to the McKinsey Global Institute, 60% of jobs could see more than 30% of their critical tasks automated, affecting 400 million to 800 million jobs by 2030.
- Remote working and e-learning will bring many opportunities to be explored by CSPs. In the short term, there are enhancements to the upstream capacity congested by the use of collaborative tools like Zoom, WebEx, and Microsoft Teams. In the long term, work, schools, and colleges will host a mixture of local and remote presence and the use of extended reality for an enhanced experience.

---

<sup>14</sup> Jordan, G., Sarmiento, T. and Campos Rodriguez, K., 2020. COVID-19 Could Cause Latam Telecommunications Revenues To Drop 3% To 14% In 2020. S&P Global - Market Intelligence.

<sup>15</sup> “The future is not what it used to be: Thoughts on the shape of the next normal”, Kevin Sneader and Shubham Singhal, McKinsey & Company, April 2020

## 5. Post COVID-19 world

### 5.1. CSPs Additional Opportunities

Aside from the COVID-19 crisis, a few CSPs are targeting new B2B and business to business to consumer (B2B2C) services to the industries which will require low jitter, latency, and strict service level agreements (SLAs) only made possible by private wireless networks coupled with edge cloud. Currently, webscale and hyperscale companies need a more granular footprint to be able to fulfill these SLAs because it is not possible with their existing centralized architectures, and here is where the battle for edge space starts - Figure 13 and Figure.

Additionally, CSPs do not own the foundational capabilities that hyperscalers have, such as agility, flexibility, and shorter innovation cycles that allow them to offer these products and services at the pace and dynamism that the clients and the market are demanding.

Current hyperscalers are approaching CSPs worldwide, given its proximity to their end clients, proposing strategy agreements to provide their customers with new services soon. Just as a final remark, given the differences described in sections 3.1 and 3.3 of the present paper, authors suggest these agreements have to be carefully analyzed and be a critical part of CSPs' business strategy.

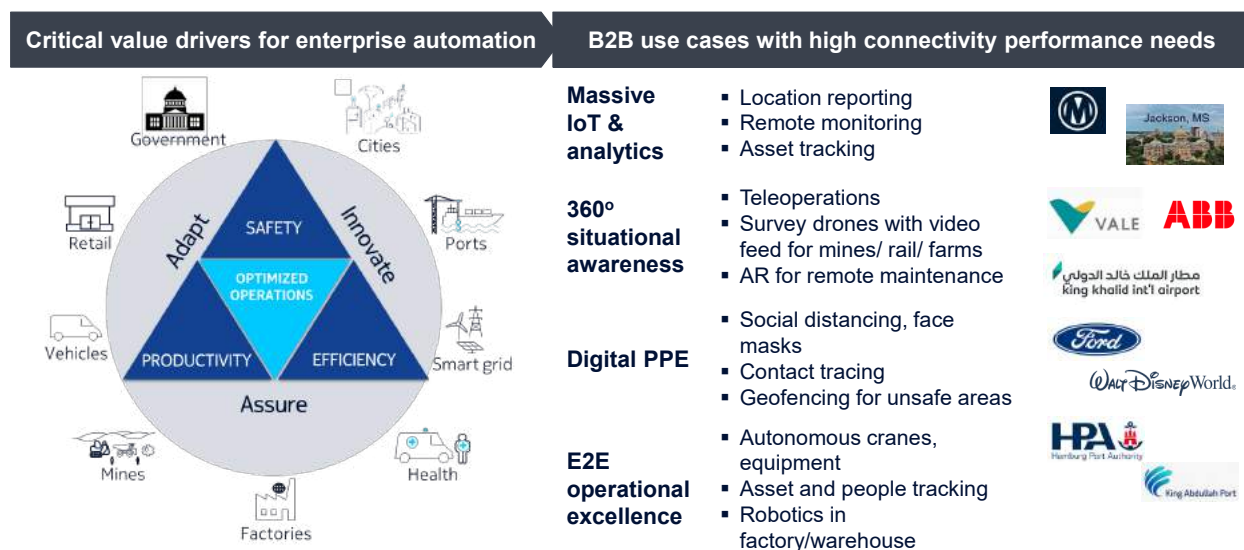


Figure 13 – Some B2B/B2B2C use cases with strict SLA – Bell Labs Consulting

## Global Server Placement Estimate




Figure 16: Providers trusted by enterprises for public edge cloud services, by vertical, 2019


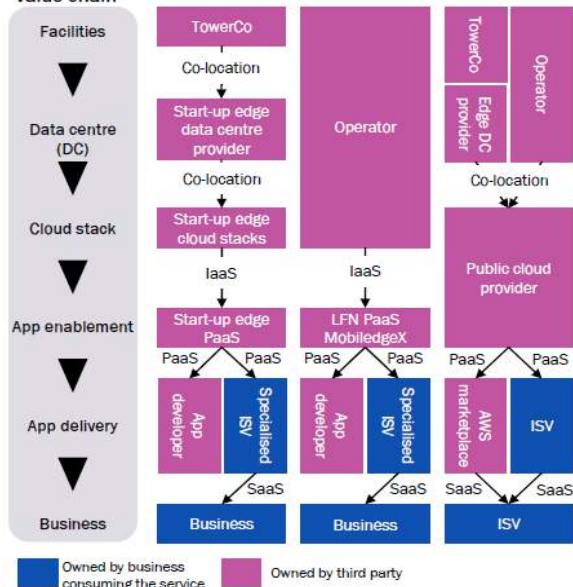


Figure 12: Three approaches to building a public edge cloud value chain




Source – "Operator opportunities and threats in the public edge cloud computing market", Analysys Mason, April 2020

### Figure 14 – Edge Battle – Footprint granularity can be a CSPs valuable asset coupled with the right strategy

To make the context more complex, a few hyperscale companies are acquiring networking vendors<sup>16</sup> and also investing in CSPs, suggesting a strong ambition to move further into the CSP territory to address high-performance services.

Since 2018, many CSPs and technology companies are experimenting with 5G use cases for the enterprise, exploring the early days of Industry 4.0 stringent SLAs and services. Today, most industries use cases are in the fields of manufacturing, transportation & logistics, energy & mining, and utilities.

<sup>16</sup> <https://blogs.microsoft.com/blog/2020/03/26/microsoft-announces-agreement-to-acquire-affirmed-networks-to-deliver-new-opportunities-for-a-global-5g-ecosystem/>



**Figure 15 - 5G announced activities by industry OMDIA July 2020**

In Table 1, there are a few public examples of CSPs and other companies using 5G access and edge cloud in gaming, public safety, and manufacturing space.

**Table 1 - 5G public announcements for enterprise - OMDIA, July 2020**

| Organization                | Type                         | Country       | Activity                   | Description                                                                                                                                                                                                                                                                                                                                                                 | Telecoms operator | 5G  | Edge computing | Date         |
|-----------------------------|------------------------------|---------------|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|-----|----------------|--------------|
| Tencent 腾讯                  | Technology vendor            | China         | Trial/test                 | Test project on Tencent's campus using a network from Huawei and China Mobile Chengdu to support MEC-hosted online gaming with low latency. (October 2019)                                                                                                                                                                                                                  | China Mobile      | Yes | Yes            | October-19   |
| OMRON INDUSTRIAL AUTOMATION | Enterprise                   | Japan         | Trial/test                 | Tripartite test program with manufacturer Omron, Nokia, and NTT DoCoMo, evaluating layout-free factory designs at multiple Omron sites. Tests include support for autonomous robots, video analytics (used to assess the differences between the most-skilled workers and everyone else), edge computing, and radio propagation in the industrial setting. (September 2019) | NTT DoCoMo        | Yes | Yes            | September-19 |
| at&t                        | Telecom operator             | United States | Trial/test                 | AT&T cooperating with Israeli drone detection startup Vorpel. Vorpel needed low latency networking and edge computing, and joined the 5G aspect of AT&T's Foundry innovation lab. (July 2019)                                                                                                                                                                               | AT&T              | Yes | Yes            | July-19      |
| at&t                        | Telecom operator             | United States | Alliance/sales partnership | Extensive partnership between AT&T and Microsoft on 5G and edge computing, including providing AI-based translation to emergency responders. (July 2019)                                                                                                                                                                                                                    | AT&T              | Yes | Yes            | July-19      |
| at&t                        | Telecom operator             | United States | Network rollout            | AT&T is collaborating with MxD (Manufacturingtimes Digital, a manufacturing industry group) to install 5G technology and Multiaccess Edge Compute (MEC) within MxD's Chicago-based innovation center. (July 2019)                                                                                                                                                           | AT&T              | Yes | Yes            | July-19      |
| BOSCH                       | Specialist/industrial vendor | Germany       | R&D project                | Bosch announces its vision for 5G-enabled future factories. Bosch aims to eliminate cabling in order to reconfigure production processes rapidly, perhaps even using wireless power. (April 2019)                                                                                                                                                                           | n/a               | Yes | Yes            | April-19     |
| Telefonica                  | Telecom operator             | Spain         | Alliance/sales partnership | Telefonica provides private network connectivity using a range of access technologies, cloud computing, and security solutions for industrial specialist Geprom's smart factory projects in Spain. Geprom provides its LEGATO SAPIENT Manufacturing Execution System management software. (April 2019)                                                                      | Telefonica        | Yes | Yes            | April-19     |

## 5.2. Three probable hypothesis

From the many outcomes of the COVID-19 impact for operators, the authors observe three probable hypotheses for CSPs' digital transformation efforts in the next few years, in Figure 16.



**Figure 16 - Digital Transformation's hypothesis after COVID-19**

In a best-case scenario, the CSP will accelerate its digital transformation efforts and will agree upon its objectives in terms of strategy definition with a holistic vision, common objective setting between top-level management, the possibility to move to different verticals, mixed and more complex value chains for new services, develop platform businesses, cannibalize current own business for a medium-long term transformation and setting levels of investment and shorter return of digital investments in its transformation plans, aiming for cost efficiencies, and additional revenues.

Besides the several moving parts, the plan would need to be executed flawlessly, with an agile organization that would follow closed-feedback loops and would be able to re-calibrate its actions to meet their main strategy. A wide range of partnership management capabilities would be instituted, and the teams would need to be able to embrace most of the challenges and successfully navigate through them.

In the second scenario, the digital transformation will be partially accomplished as some breakthroughs will be met, but the CSP may face vendor lock-in, a wide range of legacy and costly-to-replace or integrate networks and systems, which will delay their time to market and the digital transformation goals.

Last, some CSPs will start their initiatives with some guiding principles but will fall into difficult internal disputes for resources and budgets which will affect its execution. Quick wins will be achieved, for example, SD-WAN projects, but no meaningful impact on the expected cost efficiencies or in the top-line revenues.

Unfortunately, given the nature of the telecom business and its risk-averse propensity, the authors believe most CSPs will fall into the second or third categories.



### 5.3. CSPs Decisions

At this point, it is necessary to make a caveat and clarification. CSPs are not uniform entities, conversely, they constitute an innumerable amount of individual cases and particular situations. This myriad is determined by differences in some of the following aspects: strategic definitions and business volumes, contexts of regions, countries, markets, competition, technological and partnership strategies, skills, processes, and internal cultures to name a few. A few examples are the different consumption models of open-source tools<sup>17</sup>, level of disaggregation for different DSPs<sup>18</sup> or their technological partnerships, and the activities they involve.

Digital transformation outcomes will depend on each CSP and how conservative or aggressive they move in the process. The COVID-19 crisis outcome is challenging, but the fewer CSPs invest in the long-term scenario more are the chances to be disrupted by hyperscale companies.

Additionally, CSP's leadership teams should define a clear business strategy with a structured roadmap, instill a sense of urgency in the management team, and re-orient their workforce, networks, process, and tools to embrace its DSP transformation.

Probably, today more than ever, decisions will be taken with courage and an optimistic conviction that the transformation from CSP to DSP is required.

*“Everywhere we see a motley confusion which draws us into its interests, and when one thing disappears, another at once takes its place”*

*Lectures on the Philosophy of World History - Georg Wilhelm Friedrich Hegel*

## 6. Conclusion

As mentioned exhaustively, CSPs before COVID-19, had to face their DSP transformation to increase their revenues, decrease their operational expenses, provide a differentiated competitive advantage, and raise their value to its shareholders.

The pandemic was a type of black swan that can and will surely have a significant impact on the defined strategy and goals, since it affects considerably the global economy and, therefore, CSPs' investments. On the other hand, it generates opportunities to accelerate the digitalization process avoiding losing these opportunities.

So far, CSPs were successfully coping with the demands related to increasing traffic consumption, shifts in traffic patterns, upstream bottlenecks, and rural coverage shortage.

Some CSPs have been even more creative and launched interesting new services to enhance home connectivity with redundant connections, coupling broadband and wireless connectivity in the same CPE<sup>19</sup>, others were quick to attend the hospital's demand for connectivity, and some utilized IoT

---

<sup>17</sup> <https://www.lfnetworking.org/publications/2020/06/17/onap-consumption-models-whitepaper/>  
[https://www.lfnetworking.org/wp-content/uploads/sites/55/2020/06/ONAP\\_EUAG\\_Whitepaper\\_061720.pdf](https://www.lfnetworking.org/wp-content/uploads/sites/55/2020/06/ONAP_EUAG_Whitepaper_061720.pdf)

<sup>18</sup> <https://www.linkedin.com/pulse/what-your-disaggregation-tolerance-level-dtl-csp-dsp-rajiv-papneja/?articleId=6678402852722876416>  
<https://www.linkedin.com/pulse/what-your-dtl-csp-rajiv-papneja/?trackingId=2xZDItd%2FcYH3A4WzfBVR9w%3D%3D>

<sup>19</sup> <https://techround.co.uk/news/bt-launches-new-dedicated-connection-home-broadband-service/>

solutions for tracking and even experimented with drone deliveries<sup>20</sup>. Many established banking institutes sped up digital banking best practices, and notably over the top (OTT) companies have launched mobile payment solutions to support their communities<sup>21</sup>. Other companies and cities deployed temperature screening cameras to help communities to monitor fever outbreaks<sup>22</sup>. An important CSPs with a wide presence in LATAM had created a crowdsourcing app to match unemployed workers to job demands<sup>23</sup>. These initiatives are spot-on with the near-term focus on current CSPs' B2B2C competence areas.

CSPs now need to look at what they will do in the post-COVID-19 world. Possible products and services' scenario can see the acceleration of such offers, especially the ones with faster return of digital investment (RODI) and slow others given the constrained investment priorities and supply-chain shortages.

Networks will demand more upstream capacity to manage collaboration tools based on unicast implementation. The access networks are expensive to deploy and redesign and doing it to attend greater speeds or resiliency will drive a considerable part of CSP's investments. Fixed or converged operators must have a strategy that will look not only for the near term demands but beyond as the fiber return on investment (ROI) ranges for the next 7 to 10 years. Considering what-if scenarios, for example, if a new catastrophe takes place in 3 years from now, how would one network behave in terms of capacity, upstream and resiliency?

Also, given its cost-saving nature, the cloud has been deeply embraced by the IT world and slowly, we see cloud-native network applications becoming mature. CSPs are now discussing if and how they should leverage the public cloud capabilities, including edge computing capabilities, services, and placement.

The cases listed in this paper involve decisions that have to take place now, and if not well-strategized, they will have a long-term network effect on the CSP's business plans.

But the mother of all battles is still the same as CSPs had previous to this crisis. The transformation from CSPs to DSPs, as authors explain in section 3.3. Namely, this means being capable of responding to the changing demands of their clients, fostering and being part of larger ecosystems based on much more complex business models, offering not only connectivity but a wide spectrum of digital products and agile services, perhaps thinking beyond the borders of their own networks and infrastructure, and therefore increasing the value for their end customers, partners and shareholders.

This paper also depicts that an effective approach to transformation requires CSP's leadership teams to define a business vision and strategy, aligning the management organization, assess the present, and guide the transformation to the target state with a clear execution roadmap, re-orienting their workforce, networks, process, and tools to embrace the transformation program. This means that CSPs will need to have aligned the four pillars shown in Figure 5, being the business strategy, the key structure over which the other pillars must rest and always avoid the common misconception that if a company deploys an extra set of technologies, then digital transformation is done.

The pandemic crisis can force a sense of urgency to define and adopt the long-term business strategy and take advantage of the short-term opportunities, as proposed by Kotter. And doing it in the right way,

---

<sup>20</sup> <https://www.prnewswire.com/news-releases/nations-first-emergency-drone-operation-for-hospitals-pandemic-response-launches-301065751.html>

<sup>21</sup> <https://techcrunch.com/2020/06/15/whatsapp-finally-launches-payments-starting-in-brazil/>

<sup>22</sup> <https://www.reuters.com/article/us-health-coronavirus-amazon-com-cameras/exclusive-amazon-deploys-thermal-cameras-at-warehouses-to-scan-for-fevers-faster-idUSKBN2200HT>

<sup>23</sup> <http://claro.pushdobem.com.br/novo-fluxo.html>

without confusing operational efficiencies with the new business models that CSPs must generate to become DSPs.

Opportunities and the challenges are closely related and affected by these particularities such as region/country macroeconomics, strategy definitions, business models and scale, markets and competition, technologies, partnership strategies, transformation maturity, and internal cultures. Therefore, the transformation outcomes will depend on each CSP and how conservative or aggressive they move in the process, despite and exceeding the current crisis. As per the authors' experience, only a few CSPs design for a long-term strategy, leaving them prone to disruption from competitors. There are no magic recipes, short-cut solutions, or "silver bullets". Decisions will be needed with courage and an optimistic conviction that the transformation from CSP to DSP is possible.



*The challenge is enormous and never-ending, like a “LONG AND WINDING ROAD”*



# Abbreviations

|           |                                                                                      |
|-----------|--------------------------------------------------------------------------------------|
| AI        | artificial intelligence                                                              |
| BSS       | business support systems                                                             |
| B2C       | business to consumer                                                                 |
| B2B       | business to business                                                                 |
| B2B2C     | business to business to consumer                                                     |
| CPE       | customer premises equipment                                                          |
| CSP       | customer service provider                                                            |
| DSP       | digital service provider                                                             |
| CAPEX     | capital expenditure                                                                  |
| EV/EBITDA | enterprise valuation/earnings before interest, taxes, depreciation, and amortization |
| GDP       | gross domestic product                                                               |
| ICT       | information and communications technology                                            |
| IMF       | International Monetary Fund                                                          |
| ISBE      | International Society of Broadband Experts                                           |
| IT        | Information technology                                                               |
| LATAM     | Latin America                                                                        |
| ML        | machine learning                                                                     |
| NFV       | network function virtualization                                                      |
| NPS       | Net Promoter Score                                                                   |
| OPEX      | operational expenditure                                                              |
| OSS       | operations support systems                                                           |
| OTT       | over the top                                                                         |
| RAN       | radio access network                                                                 |
| RGW       | residential gateway                                                                  |
| RODI      | return of digital investment                                                         |
| ROC       | return on capital                                                                    |
| ROI       | return of investment                                                                 |
| SCTE      | Society of Cable Telecommunications Engineers                                        |
| SDN       | software-defined networks                                                            |
| SLA       | service level agreement                                                              |
| TCO       | total cost of ownership                                                              |
| TTM       | time to market                                                                       |
| VPN       | virtual private network                                                              |

## Bibliography & References

<https://www.weforum.org/focus/fourth-industrial-revolution>

<https://es.wikipedia.org/wiki/Neuralink>

<https://www.forbes.com/sites/cathyhack1/2020/07/05/the-metaverse-is-coming--its-a-very-big-deal/#3cc83452440f>

[https://en.wikipedia.org/wiki/Mirror\\_world](https://en.wikipedia.org/wiki/Mirror_world)

<https://www.mckinsey.com/business-functions/mckinsey-digital/our-insights/the-trillion-dollar-opportunity-for-the-industrial-sector>

Creaner, M., 2020. Transforming The Telco: Developing The Instinct For Survival In The Digital Economy. CenterNODE.org. <https://www.youtube.com/watch?v=xl6gIEb8OAM> (Martin Creaner - Former CEO & President - TM Forum)

“Digital Transformation Is About Talent, Not Technology”, Becky Frankiewicz and Tomas Chamorro-Premuzic, Harvard Business Review, May 2020. Available online at: <<https://hbr.org/2020/05/digital-transformation-is-about-talent-not-technology>>

"Leading Change - Why Transformation Efforts Fail", John P. Kotter, Harvard Business Review, January 1996

“Trash the Rulebook”, Accenture, 2018

“The Two Big Reasons That Digital Transformations Fail”, Mike Sutcliffe, Raghav Narsalay and Aarohi Sen, Harvard Business Review, October 2019. Available online at: <<https://hbr.org/2019/10/the-two-big-reasons-that-digital-transformations-fail>>

“6 Digital Strategies, and Why Some Work Better than Others”, Jacques Bughin and Nicholas Van Zeebroeck, Harvard Business Review, July 2017

<https://www.imf.org/en/Publications/WEO/Issues/2020/06/24/WEOUpdateJune2020>

<https://www.worldbank.org/en/publication/global-economic-prospects>

"The Global Internet Phenomena Report COVID-19 Spotlight", Sandvine, May 2020

“Acciones frente al COVID-19”, Telecom Argentina, April 2020

“The future is not what it used to be: Thoughts on the shape of the next normal”, Kevin Sneader and Shubham Singhal, McKinsey & Company, April 2020

"Operator opportunities and threats in the public edge cloud computing market“, Analysys Mason, April 2020

“CSPs Must Make Crafting Their Digital Dragon Strategy a Top Priority”, Gartner, May 2020

"Edge Computing: from standard to actual infrastructure deployment and software development", INTEL, October 2019

“Top 10 Strategic Technology Trends for 2020: Empowered Edge”, Gartner, March 2020

“Market Guide for Edge Computing Solutions for Industrial IoT”, Gartner, September 2019

<https://blogs.microsoft.com/blog/2020/03/26/microsoft-announces-agreement-to-acquire-affirmed-networks-to-deliver-new-opportunities-for-a-global-5g-ecosystem/>

<https://www.lfnetworking.org/publications/2020/06/17/onap-consumption-models-whitepaper/>

[https://www.lfnetworking.org/wp-content/uploads/sites/55/2020/06/ONAP\\_EUAG\\_Whitepaper\\_061720.pdf](https://www.lfnetworking.org/wp-content/uploads/sites/55/2020/06/ONAP_EUAG_Whitepaper_061720.pdf)

<https://www.linkedin.com/pulse/what-your-disaggregation-tolerance-level-dtl-csp-dsp-rajiv-papneja/?articleId=6678402852722876416>

<https://www.linkedin.com/pulse/what-your-dtl-csp-rajiv-papneja/?trackingId=2xZDItd%2FcYH3A4WzfBVR9w%3D%3D>

<https://techround.co.uk/news/bt-launches-new-dedicated-connection-home-broadband-service/>

<https://www.prnewswire.com/news-releases/nations-first-emergency-drone-operation-for-hospitals-pandemic-response-launches-301065751.html>

<https://techcrunch.com/2020/06/15/whatsapp-finally-launches-payments-starting-in-brazil/>

<https://www.reuters.com/article/us-health-coronavirus-amazon-com-cameras/exclusive-amazon-deploys-thermal-cameras-at-warehouses-to-scan-for-fevers-faster-idUSKBN2200HT>

<http://claro.pushdobem.com.br/novo-fluxo.html>

<https://www.linkedin.com/pulse/covid-19-pandemic-catalyst-4th-industrial-revolution-changes-ger/>

COVID-19 creates pain, change and even pockets of opportunity for the IT Industry, TBR Special Report, March 20, 2020

# Field Experiences with US OFDMA and using US Profile Management

A Technical Paper prepared for SCTE•ISBE by

**Karthik Sundaresan**

Distinguished Technologist  
CableLabs

858 Coal Creek Circle, Louisville, CO 80303  
3036613895  
k.sundaresan@cablelabs.com

**João Pedro Fernandes**

HFC Engineer  
NOS

Portugal  
+ 351 931005187  
joao.fernandes@nos.pt

**Jay Zhu**

Senior Engineer  
CableLabs

858 Coal Creek Circle, Louisville, CO 80303  
3036613312  
j.zhu@cablelabs.com

# Table of Contents

| Title                                                                         | Page Number |
|-------------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                          | 4           |
| 2. US OFDMA background .....                                                  | 4           |
| 2.1. US splits .....                                                          | 4           |
| 2.2. OFDMA channel basics .....                                               | 4           |
| 2.3. US OFDMA Frames & minislots.....                                         | 5           |
| 2.4. US FEC .....                                                             | 6           |
| 2.5. Upstream Noise funneling .....                                           | 6           |
| 2.6. Common Upstream Issues.....                                              | 7           |
| 2.7. CMTS Profile / IUC mangement.....                                        | 8           |
| 2.7.1. Upstream Profile Testing .....                                         | 8           |
| 2.7.2. Upstream Probes and RxMER Measurements.....                            | 8           |
| 2.7.3. Upstream Data Profile Testing Bursts .....                             | 8           |
| 2.7.4. IUC/ Profile change .....                                              | 9           |
| 2.8. Upstream OFDMA : the need for Profile Management .....                   | 9           |
| 2.9. Designing Profiles for US OFDMA channel .....                            | 10          |
| 3. US OFDMA Deployment at NOS .....                                           | 11          |
| 3.1. Channel location .....                                                   | 11          |
| 3.2. Phases of turning on OFDMA .....                                         | 12          |
| 3.3. Technology maturity.....                                                 | 13          |
| 3.4. Initial IUC definition for OFDMA chanel .....                            | 13          |
| 3.4.1. Data collection for initial IUC.....                                   | 14          |
| 3.4.2. Initial IUC Definition .....                                           | 14          |
| 3.5. Upstream Plant readiness.....                                            | 15          |
| 3.6. KPIs for monitoring.....                                                 | 15          |
| 3.7. US RxMER Data Collection from the CMTS.....                              | 17          |
| 3.7.1. Standard SNMP + TFTP method .....                                      | 17          |
| 3.7.2. SNMP + SFTP method .....                                               | 17          |
| 3.7.3. CLI method.....                                                        | 17          |
| 3.8. US RxMER Data Collection in field deployment for PMA trials .....        | 18          |
| 4. US PMA Algorithms.....                                                     | 18          |
| 4.1. Differences from Downstream Algorithms .....                             | 19          |
| 4.2. Percentile method .....                                                  | 19          |
| 4.2.1. Algorithm 1A : Per CM Percentile .....                                 | 19          |
| 4.2.2. Algorithm 1B: All CMs Percentile .....                                 | 20          |
| 4.2.3. Algorithm 1C: Remove Outliers + Percentile .....                       | 20          |
| 4.3. Time Clustering Methods .....                                            | 21          |
| 4.3.1. Algorithm 2A: Time Clustering using Actual CM Samples.....             | 21          |
| 4.3.2. Algorithm 2B: Time clustering using artificial Percentile Samples..... | 21          |
| 4.3.3. Algorithm 2C: Time Clustering after removing outliers .....            | 22          |
| 4.4. Other PMA Considerations .....                                           | 22          |
| 5. US PMA Field Trial Results.....                                            | 23          |
| 5.1. US RxMER Field Data .....                                                | 23          |
| 5.2. US IUCs/Profiles designed.....                                           | 23          |
| 5.3. Performance and Stability .....                                          | 26          |
| 5.4. US IUCs/Profiles designed with no IUC limitations.....                   | 26          |
| 6. Conclusion / Future Work.....                                              | 27          |
| Abbreviations .....                                                           | 28          |
| Bibliography & References.....                                                | 28          |

## List of Figures

| Title                                                                                                                     | Page Number |
|---------------------------------------------------------------------------------------------------------------------------|-------------|
| Figure 1 – Upstream OFDMA Minislot Layout and Grants across Minislots .....                                               | 5           |
| Figure 2 –Sample Pilot Patterns .....                                                                                     | 6           |
| Figure 3 – US Noise funneling .....                                                                                       | 7           |
| Figure 4 – US RxMER Measured on multiple CMs on a upstream Channel.....                                                   | 10          |
| Figure 5 – Target spectrum for OFDMA channel placement.....                                                               | 11          |
| Figure 6 - Example power level data used to calculate initial IUC definitions .....                                       | 14          |
| Figure 7 - Example IUC definition from the deduced MER.....                                                               | 14          |
| Figure 8 – Tracking IUC-usage-hours per day (across CMTS) .....                                                           | 16          |
| Figure 9 – Operational dashboards displaying CM IUC distribution in one US SG and top US SGs with<br>CM impairments ..... | 16          |
| Figure 10 – Us RxMER Samples from 4 CMs .....                                                                             | 23          |
| Figure 11 – IUCs created using Percentile Algorithms 1A,1B,1C .....                                                       | 24          |
| Figure 12 – IUCs created using Time Clustering Algorithms 2A, 2B, 2C .....                                                | 25          |
| Figure 13 – 7 IUCs created using all Algorithms .....                                                                     | 27          |

## List of Tables

| Title                                                         | Page Number |
|---------------------------------------------------------------|-------------|
| Table 1 – D3.1 Upstream OFDMA Parameters .....                | 5           |
| Table 2 - FEC coding parameters .....                         | 6           |
| Table 3 – Upstream Channel Parameters Chosen on CMTSs.....    | 12          |
| Table 4 – Field Results from using PMA Algorithm 1A, 1B ..... | 26          |

# 1. Introduction

DOCSIS® 3.1 is now largely deployed in the field, but so far it has been an Orthogonal Frequency Division Multiplexing (OFDM) Downstream only endeavor. Operators are beginning to test Upstream Orthogonal Frequency Division Multiple Access (OFDMA) and are discovering various intricacies in getting the US OFDMA to work robustly. We CableLabs and NOS, a cable operator in Portugal have been working together and have been focused on DOCSIS 3.1 OFDMA in general and collaborating on the Upstream Profile Management Application (PMA). Defining appropriate Upstream profiles for an OFDMA channel affects the stability of the modems on the channel and also the capacity realized. We implemented data collection agents through the CMTS CLI to collect Upstream RxMER for each CM, or on other CMTSs collect the data manually for each CM. US RxMER looks very different than the Downstream RxMER, due to the noise funneling characteristics on the HFC plant. NOS engineering and CableLabs collaborated on various US PMA algorithms and this paper will describe some of the new methodologies we have developed. The NOS operations team is doing a field trial with the US data collection and then configuring the profiles/IUC (interval usage code) generated by PMA on a live plant to understand the impact and behavior. This paper focuses on the lessons we have learnt from the DOCSIS 3.1 upstream field trial and production systems.

## 2. US OFDMA background

DOCSIS 3.1 introduces OFDM downstream signals and OFDMA upstream signals to achieve robust operation and provide more efficient use of the spectrum than previous DOCSIS versions. OFDMA for the upstream path is a multi-user version of OFDM, and assigns subsets of subcarriers to individual CMs.

### 2.1. US splits

The DOCSIS 3.1 system will have options of several split configurations that can be exercised based on traffic demand, services offered and the capability of the cable plant. In the upstream direction, the cable system may have a 5-42 MHz, 5-65 MHz, 5-85 MHz, or 5-204 MHz pass bands. A D3.1 CM supports one or more of the following upstream upper band edges, (as long as one is 85 MHz or greater): 42 MHz; 65 MHz, 85 MHz, and/or 204 MHz. The DOCSIS 3.1 Network supports a minimum of two independently configurable OFDMA upstream channels with each occupying a spectrum of up to 95 MHz. A DOCSIS 3.1 CM is capable of transmitting on OFDMA channels and legacy single carrier-QAM channels (SC-QAM) at the same time (as controlled by the CMTS). There are no legacy SC-QAM channels above a frequency of 85 MHz.

### 2.2. OFDMA channel basics

The OFDMA upstream multicarrier system is composed of either 25 kHz or 50 kHz wide subcarriers. In the upstream, the subcarriers are grouped into independently configurable OFDMA channels each of up to 95 MHz encompassed spectrum, totaling 3800 25 kHz spaced subcarriers or 1900 50 kHz spaced subcarriers. When configured for 2K FFT (Fast Fourier Transform), the CMTS uses the subcarriers in the range  $74 \leq k \leq 1973$ , where  $k$  is the index of the subcarrier defining the OFDMA signal. When configured for 4K FFT, the CMTS number uses subcarriers numbered in the range  $148 \leq k \leq 3947$ .

The parameters of the two OFDMA channels can be independently configured thereby optimizing configuration based on channel conditions. The table lists the Upstream Channel parameters, from [PHYv3.1]



**Table 1 – D3.1 Upstream OFDMA Parameters**

| Parameter                           | Value                                                                                                                                                                   |                      |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| Upstream Sampling Rate              | 102.4 MHz                                                                                                                                                               |                      |
| Upstream Elementary Period Rate     | 1/102.4 MHz                                                                                                                                                             |                      |
| Channel bandwidth                   | 10 MHz, ..., 95 MHz                                                                                                                                                     | 6.4 MHz, ..., 95 MHz |
| IDFT size                           | 2048                                                                                                                                                                    | 4096                 |
| Subcarrier spacing                  | 50 kHz                                                                                                                                                                  | 25 kHz               |
| Symbol duration                     | 20 $\mu$ s                                                                                                                                                              | 40 $\mu$ s           |
| Maximum active subcarriers (95 MHz) | 1900                                                                                                                                                                    | 3800                 |
| OFDMA Cyclic Prefix size            | 0.9375 $\mu$ s, 1.25 $\mu$ s , 1.5625 $\mu$ s , 1.875 $\mu$ s<br>2.1875 $\mu$ s , 2.5 $\mu$ s, 2.8125 $\mu$ s<br>3.125 $\mu$ s, 3.75 $\mu$ s, 5.0 $\mu$ s, 6.25 $\mu$ s |                      |
| OFDMA Roll-off Period Size          | 0 $\mu$ s, 0.3125 $\mu$ s , 0.625 $\mu$ s , 0.9375 $\mu$ s<br>1.25 $\mu$ s , 1.5625 $\mu$ s, 1.875 $\mu$ s, 2.1875 $\mu$ s                                              |                      |
| OFDMA Modulation orders             | (BPSK), QPSK, 8-QAM, 16-QAM, 32-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, 1024-QAM, 2048-QAM, and 4096-QAM                                                                |                      |

### 2.3. US OFDMA Frames & minislots

DOCSIS 3.1 Upstream transmission uses OFDMA frames. Each OFDMA frame is comprised of a configurable number of symbols ( $K = 6$  to 36). Several transmitters may share the same OFDMA frame by transmitting on allocated subcarriers of the OFDMA frame. The structure of an OFDMA frame is depicted in Figure 1. The upstream spectrum is divided into groups of subcarriers called minislots. Minislots have dedicated subcarriers, all with the same modulation order ('bit loading'). [PHYv3.1] specifies two minislot sizes by specifying the number of subcarriers per minislot. There are 8 or 16 subcarrier minislots, a minislot is always 400KHz wide (25KHz subcarrier \*16, or 50KHz subcarrier \*8). Minislots have dedicated subcarriers, all with the same modulation order ('bit loading'). Though the span of the minislot is always 400KHz, the length of the minislot in the time is the same as the number of symbols( $K$ ) of the frame. An operator can configure the number of symbols in an OFDMA frame to pick an appropriate size for the minislots on a channel. A CM is allocated to transmit one or more minislots in a transmission Burst. The modulation order of a minislot, as well as the pilot pattern used may change between different transmission bursts and are determined by the profile definition. Several transmitters may share the same OFDMA frame by transmitting on their allocated minislots on the OFDMA frame.




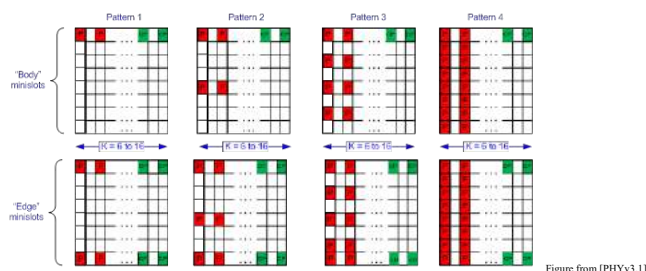
Figure from [MULPh3.1]

**Figure 1 – Upstream OFDMA Minislot Layout and Grants across Minislots**

There are two types of minislots edge minislots and body minislots. An edge minislot is the first minislot in a transmission burst, and body minislots are used for all other minislots in a transmission burst. See [PHYv3.1] for minislot usage with exclusion bands etc.

Each minislot is comprised of pilots (P), complementary pilots (CP), and data subcarriers. Pilots are used by the CMTS receiver to adapt to channel conditions and frequency offset. Pilots are subcarriers that do not carry data and encode a pre-defined BPSK symbol known to the receiver. [PHYv3.1] also specifies complementary pilots which are subcarriers that carry data, but with a lower modulation order than other data subcarriers in the minislot. If the modulation order used for data subcarriers in the minislot is M, the complementary pilots are used with modulation order equal to the maximum between M-4 and 1 (BPSK). For example, if the bit loading in a minislot is 10, Complementary Pilots use 6 bits.

For each minislot size, seven pilot patterns are defined, see figure 2. Pilot patterns differ by the number of pilots in a minislot, and by their arrangement within the minislot. The different pilot patterns enable the operator to optimize its performance (physical layer rate and pilot overhead) according to different conditions and variations of SNR with frequency. Each pilot pattern defines edge and body minislots.



**Figure 2 – Sample Pilot Patterns**

## 2.4. US FEC

An upstream grant from the CMTS indicates which minislots are assigned to a given burst and which upstream profile is to be used by the CM. The CM and CMTS use this information to determine the total number of bits in the grant which are available to be used for FEC information or parity. Per [PHYv3.1], OFDMA use three Quasi-Cyclic Low-Density Parity-Check codes (QC-LDPC) for the upstream transmission, as depicted in Table below

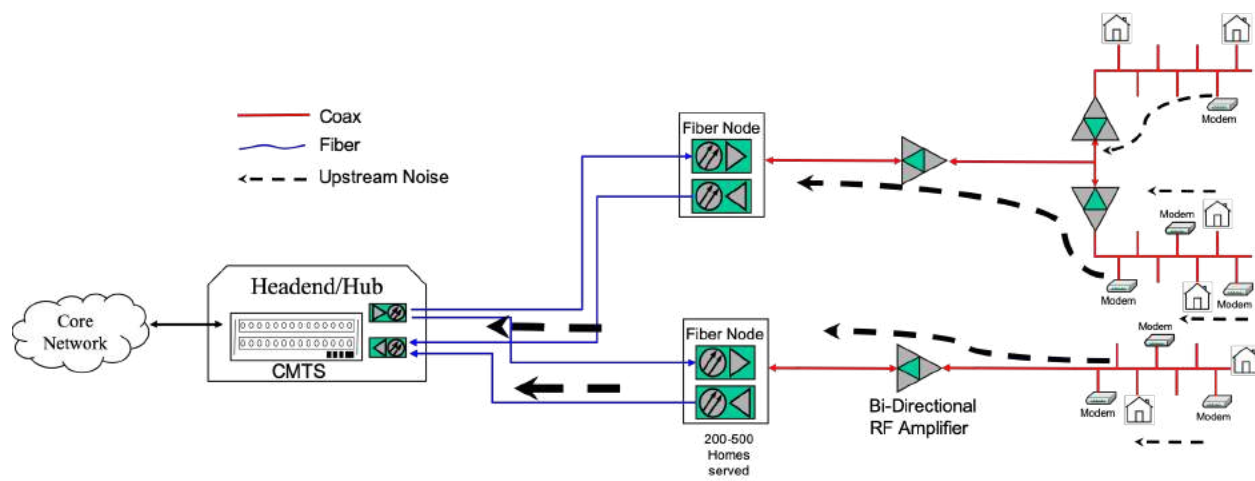
**Table 2 - FEC coding parameters**

| Code   | LDPC Code Rate | Codeword size in bits | Information bits | Parity bits |
|--------|----------------|-----------------------|------------------|-------------|
| Long   | $8/9 = 89\%$   | 16200                 | 14400            | 1800        |
| Medium | $28/33 = 85\%$ | 5940                  | 5040             | 900         |
| Short  | $3/4 = 75\%$   | 1120                  | 840              | 280         |

## 2.5. Upstream Noise funneling

The DOCSIS upstream behaves differently than the DOCSIS downstream. As the DOCSIS downstream signal is transmitted (from the CMTS to the CM) through the branched cable topology, the signal becomes attenuated and weaker at each branching and with cable distance. The strong signal at the headend now requires amplification to adequately reach the subscriber. Noise or interference in

downstream frequencies that enter the cable system say at a node or an amplifier only affects customers past that point.



**Figure 3 – US Noise funneling**

In the downstream direction, there is one location from where the signals enter the HFC plant, specifically the cable modem termination system (CMTS) in the headend or at a node location with distributed access architectures (e.g. Remote PHY). The operator has control of the signal at that point and along the network, to ensure it reaches every CM. From the headend to the CM, the RF signal fans out in a star topology network in a point to multipoint fashion. It is the opposite on the upstream/return path: the RF signals enter the plant from every home that is attached to the plant, and all of those signals combine together as they travel to the headend. Typical of all point to multipoint networks, the noise from every device on the network gets combined as it travels upstream and is finally received on the upstream port at the CMTS. This is known as the noise funneling problem, as shown in the figure 3. The additive nature of noise has a large impact at the CMTS upstream receiver. A small amount of noise from every unterminated cable or loose connector enters the upstream path and is combined with other ingress noise and then amplified as it travels upstream and gets funneled up to the CMTS causing it to be unable to decode the communication from the CM. HFC plant segmentation somewhat mitigates the noise funneling effect by reducing the number of combined signals traveling on a given upstream path.

## 2.6. Common Upstream Issues

There are many upstream impairments on the cable plant, some of the causes and ways operators tackle them are described here. Cable operators are constantly working to mitigate upstream ingress noise. Cable operators have a hard time pinpointing where ingress noise is entering the plant. A spectrum analyzer can allow an operator to visualize the noise problem, but it doesn't tell the operator the location of the problem. A common method, to track ingress, is disconnecting each segment of the plant at a node to see the impact on noise levels on the spectrum analyzer. This is laborious and customer impacting and as the noise sources are bursty and intermittent, there is a bit of luck involved in the technician finding the source. A technician needs to visit the worst ingress locations, and make physical changes to the network. Often the ingress locations are also a place where signals leak out of the cable plant. There are many tools available which are based on this concept of looking for test signals leaking out of the plant.

Thermal noise at the amplifiers and fiber optic link noise can be sources of upstream impairments. Other ingress noise sources on the upstream path include impulse noise from loose connectors, reflections from unterminated splitters or taps, cracked cables, common path distortion due to corroded connectors or

cabling issues. Many of these are resolved by physical repair to the damaged cable network infrastructure (replacing drop cable, servicing defective network components such as power supplies). Reflections can be mitigated through equalization coefficients used within the DOCSIS technology. Identifying problem frequencies and avoiding them for upstream transmission is another method, to side step some problems at least temporarily. Using external data analytics an operator can potentially identify ingress locations and use frequency changes to dynamically avoid this type of noise. Impulse noise can be partially mitigated by the use of DOCSIS forward error correction (FEC) and interleaving algorithms. Again, a technician might need to visit the worst locations to make physical changes to the network. Correct amplifier alignment and setting up correct power levels in the upstream path also prevents issues like laser clipping.

## **2.7. CMTS Profile / IUC management**

The CMTS assigns OFDMA Upstream Data Profile (OUDP) IUCs to the CMs based on the measured plant conditions. It is intended that the Data Profile IUC 13 is configured as a robust OFDMA profile usable by any DOCSIS 3.1 CM served by that upstream channel. Data Profile IUC 13 is used for all OFDMA data grants to modems which have not completed registration.

During or after modem registration, the CMTS has the option of assigning the CM to use any other configured data profile. Typically, the data profiles other than IUC 13 will be configured for higher performance than IUC 13, although not all of these profiles will be usable by all modems. The CMTS assigns the CM either one or two data profiles (IUCs) for each OFDMA channel in the modem's Transmit Channel Set. This can be assigned during Registration, and can be changed after Registration using dynamic messages (DBC transaction). After registration, the CMTS grants bandwidth on the OFDMA channel for data transmissions to a CM using one of the CM's assigned OUDP IUCs.

### ***2.7.1. Upstream Profile Testing***

Because it is expected that not all upstream data profiles will be usable by all CMs, a CMTS can evaluate a CM's performance using a particular profile before assigning that profile to be used for live user traffic. The DOCSISv3.1 technology [MULPIv3.1] provides various tools to aid the CMTS in gathering information about upstream profile performance. A CMTS performs such an evaluation in vendor-specific ways, usually revolving around modulation error ratios or codeword error ratios. These tools are based on two types of transmissions: upstream probes, and upstream Data Profile Testing bursts.

### ***2.7.2. Upstream Probes and RxMER Measurements***

A CMTS uses upstream probes for ranging-related functions such as determining transmit pre-equalizer coefficients. A CMTS also has the option of using an upstream probe to take an RxMER (received modulation error ratio) measurement. The CMTS grants probe opportunities to a CM in a P-MAP message with the "MER" bit set. When the CMTS receives the probe transmissions from the CM corresponding to such a grant, it performs the RxMER measurement and uses the results in its decision making. It also populates the corresponding MIB object or can upload a RxMER per subcarrier file via TFTP, for the operator's information.

### ***2.7.3. Upstream Data Profile Testing Bursts***

Some types of upstream profile performance cannot be measured using probe bursts. For example, a CMTS might wish to gather information on FEC performance or count CRC errors for a particular profile. Probe bursts cannot be used for these purposes as don't carry any information, and instead the CM&CMTS can send/receive upstream Data Profile Testing bursts. Per [MULPIv3.1], to command a CM

to send an upstream Data Profile Testing burst, the CMTS first assigns an OUDP Testing SID to the CM on one or more upstream channels. The CMTS then sends a grant to an OUDP Testing SID, the IUC of this grant is an Assigned OUDP IUC currently assigned to the modem. The modem responds to a valid grant to any of its OUDP Testing SIDs by sending a Data Profile Testing burst in the grant. The Data Profile Testing burst from the CM is a 64-byte Ethernet packet, with counting pattern in payload bytes beginning with 0x01, continuing with 0x02, 0x03, etc., and ending with 0x2E (count is re-started at 0x01 in each successive packet). The CM fills the grant with DOCSIS frames. The modem treats all grants to its OUDP Testing SID(s) as grants to a single flow existing across all OFDMA channels to the SID has been assigned.

#### **2.7.4. IUC/ Profile change**

The CMTS assigns one or two OUDP IUCs to a CM, once the assignment is successful, that CM is ready for transmitting data using the assigned IUCs. After registration, the CMTS grants OFDMA bandwidth for data transmissions to a CM using one of the CM's assigned OUDP IUCs.

A CM supports 2 US Profiles/IUCs at a time. A CM starts on the OFDMA channel with IUC 13 (e.g. say set to 16 QAM). At a later point the CM is assigned an additional IUC (e.g. IUC 12, say 256 QAM). When CMTS sees US FEC errors on the secondary profile (IUC 12 in this example), it chooses to rectify the situation. A CMTS can reassign the CM a new IUC, say IUC 11 (with 64 QAM in areas of high noise and 256 QAM elsewhere) dynamically via DBC messages. The CMTS continues to use the default IUC-Profile 13 to forward traffic to avoid any packet loss during IUC change, when the DBC is in process. In practice, this means that the upstream capacity for the CM is changing intermittently as it switches between profiles, which could lead to a degraded performance and user experience.

## **2.8. Upstream OFDMA : the need for Profile Management**

The DOCSIS 3.1 specification fundamentally changes the nature of information delivery on the cable plant, and the way HFC networks will be maintained and managed. In a significant change from previous DOCSIS versions, the OFDM/OFDMA channel does not use a one-size-fits-all modulation scheme; rather, the modulation can be optimized based on actual plant conditions at different frequencies and individual devices. CMs and CMTS that communicate with a cleaner signal can utilize an efficient high-order-modulation, while devices that have a degraded signal will use more robust modulation, all on the same channel.

The DOCSIS 3.1 toolbox provides a wide range of modulation choices that can be used to fine-tune the transmissions to get the best performance from the current network conditions. To manage the optimization of these settings across the population of devices, the CMTS uses the concept of Upstream profiles. An upstream profile defines the modulation order (i.e., bit loading) and pilot pattern on each of the minislots on the channel (up to 237), spanning (up to 3800 or 1900 subcarriers) across the OFDMA channel.

DOCSIS 3.1 specifications [MULPIv3.1] provides for defining multiple upstream profiles, where each profile can be tuned to account for the specific plant conditions that are experienced by a set of CMs. A well-designed, optimized set of modulation profiles allows an upstream channel to operate with robustness and a lower SNR margin, potentially allowing a channel to deliver an overall higher throughput. In addition, it can allow for communication to devices by providing service even in situations where significant plant impairments exist.

The application that implements this optimization logic is external to a CMTS, enabling the most efficient use of profiles across channels and CMs. For an operator, it also allows uniform operation of such

algorithms across different CMTS platforms. This profile optimization and profile creation functionality is implemented as an ‘application’ running outside the CMTS and is known as the Profile Management Application (PMA). Managing profiles manually for an operator for thousands of CMTS Upstream channels is a labor-intensive process requiring a deep understanding of the channel conditions. The calculation and recalculation of profiles would overload human operators and PMA simplifies this process significantly.

## 2.9. Designing Profiles for US OFDMA channel

Now in the upstream, the noise from every house and every network element gets accumulated and is seen at the upstream receiver on the CMTS. Now a CMTS receiver can measure the received modulation error ratio (RxMER) for each CM, see figure 4 for some example measurements from a live network. In the upstream, this signal to noise signature for each of the CMs (that are sharing the upstream channel) starts looking very similar, as they all share the same noise across the channel with slight differences due to the signal levels itself or some in house network problems. This means common profiles can be designed for many CMs experiencing similar noise conditions and most CMs will be able to use a common profile. (This is very different from the behavior seen in the downstream, where different sets of CMs have very different noise signatures.) The variation in the Upstream RxMER from sample to sample for a single modem itself is much greater than the mostly tight RxMER variations that we are accustomed to seeing in the downstream. For CMs which suffer more noise, they can be put into a different profile optimized for their particular noise environment. The modulation orders within a profile can vary appropriately across the spectrum as per the noise levels in that part of the spectrum.



**Figure 4 – US RxMER Measured on multiple CMs on a upstream Channel**

The upstream Profile Management Application (PMA) can automate this design of the profiles on upstream channels across various segments in the cable plant. Reading the upstream RxMER from the CMTSs on the network, processing the RxMER information with intelligent algorithms to create profiles, and then configuring the newly optimized profiles on the CMTS are the primary functions an upstream PMA solution accomplishes. Configuring optimized profiles brings solid reliability to the upstream network connection and also increases the capacity in parts of the spectrum which can accommodate higher modulation orders.

Given this understanding of the upstream plant behavior and the RxMER signatures, the question now is, what are the best algorithms to design upstream profiles which give operators robust upstream operation as well as increase the throughput? We address this problem with a few different solutions in chapter 4.



### 3. US OFDMA Deployment at NOS

In this chapter we share some of the Upstream OFDMA field rollout experiences by the NOS access engineering team in Portugal earlier this year. After some lab trials the effort quickly moved to limited field trials with internal employees and then to customers, in a phased approach. The support for OFDMA on CMTSs and CMs is maturing, but some of the initial trials led to a lot of lessons learned. Many of these bug fixes and feature improvements may take time to make it into operator networks, and we hope that sharing those notes here helps other operators with their Upstream trials and roll-outs. There may be gaps in system implementations which could use some thought and new solutions. Data collection for Upstream is based on the CMTS and these capabilities are still limited across the different implementations. This data is necessary to ensure that an operator can design the correct IUCs for each of the US OFDMA channels. Like many European HFC plants, the NOS plant has a 65 MHz EuroDOCSIS® split, which leaves open a reasonable amount of spectrum of OFDMA channels. We discuss the reasons we decided to keep the trials above 23.5 MHz and also share some trial experiences below that. The HFC upstream plant clean-up is always good practice for all operators and this trial needed more of that to make the OFDMA operations more robust. Monitoring the network and developing a few Key Performance Indicators (KPIs) for the DOCSIS Upstream are important to evaluate the changes made to the network and understand the performance of the network.

#### 3.1. Channel location

The NOS plant in Portugal has upstream spectrum up to 65 MHz and currently there are 2 to 3 SC-QAMs per serving group and up to 4 upstream serving groups (US SGs) per downstream serving groups (DS SGs). With a 65MHz split and SC-QAM channels being already located at the upper edge, the natural choice to place OFDMA was the lower part of the spectrum, with a hope to take advantage of the enhanced flexibility of OFDMA.



**Figure 5 – Target spectrum for OFDMA channel placement**

The initial trials started with OFDMA channels located at 15MHz up to 45.5 or 51.9MHz, depending on existing number of total SC-QAMs in the US SG, which were either 2 or 3. The spectrum below 40MHz was never used before for DOCSIS upstream and the initial analysis indicated that a high level of noise was sometimes present. Even after a fair bit of work to clean the outside plant, it turned out that using spectrum below 23.5MHz made the channel too susceptible to impulse noise. The result of this was that the OFDMA channel was affected, leading to channel impairment. So, the decision was made to locate the lower edge of the OFDMA channel at 23.5MHz for better stability.

At the time of this writing another approach was being tested: This approach chose to locate the SC-QAMs from 30MHz and use the upper part of the upstream spectrum for the OFDMA channel. As this is a much cleaner spectral area, more bits/hz can be extracted from the same bandwidth for the OFDMA

channel. Now this will only work if the SC-QAMs can work error-free in the lower zone of the spectrum. To ensure the stability of the US SC-QAMs, upstream agility features will definitely help.

Another approach that needs to be explored is to use two smaller OFDMA channels instead of a larger, single one. The idea is to have each of these OFDMA channels straddle the 2-3 SC-QAM channels. This way a smaller but highly reliable OFDMA channel can be obtained, providing a baseline capacity that is always present. The second OFDMA channel, due to its location in the lower end of the spectrum, will be much more susceptible to impairment or US IUC downgrades/flapping. This channel can then act as reserve capacity available most of the time to most of the CMs.

The field deployment was on two different CMTS platforms, both with different upstream capabilities, the table below describes some of the parameters chosen for the deployment.

**Table 3 – Upstream Channel Parameters Chosen on CMTSs**

| Parameter            | CMTS vendor 1        | CMTS vendor 2                      |
|----------------------|----------------------|------------------------------------|
| OFDMA Bandwidth      | 20.1 – 45.5/51.9 MHz | 15.5 - 45.5 MHz<br>21.9 - 51.9 MHz |
| Number of IUCs       | 2 IUCs – 13; 12      | 2 IUCs – 13; 12                    |
| Variable bit loading | Flat Profiles        | Variable                           |
| K                    | 18                   | 16                                 |
| Pilot Pattern        | 4                    | 2                                  |
| Subcarrier Spacing   | 50                   | 50                                 |
| Roll Off             | 224                  | 96                                 |
| Cyclic Prefix        | 320                  | 192                                |

85 MHz Trial Learnings: In limited areas we also tried out OFDMA channels on an 85 MHz plant, as that was available in some areas. Here the OFDMA channel was located from 22MHz to 85MHz – with an exclusion band for 2 SC-QAM channels. Initially for the OFDMA channel we needed to create a 1.6MHz band in each side of the QAMs with 64QAM modulation, as per vendor recommendations, to protect the rest of the OFDMA channels. This configuration was revealed to be unstable with loss of IUC12 and OFDMA impairment. What solved the problems was adding a 500kHz guard band at the exclusion zone. Once this was added and the Upstream performed well. The calculated OFDMA capacity was 344 Mbps, and along with the couple of SC-QAM channels, actual speed tests gave us numbers of ~400 Mbps. We plan to use this selectively in very noisy and tough upstream environments, reusing existing 85MHz-ready equipment when available.

### 3.2. Phases of turning on OFDMA

OFDMA is a new technology for the Operators and for the CMTS and CM vendors, so a very cautious approach was used to deploy this technology in the network and enabling traffic on those OFDMA channels. After some initial lab trials, a three-phase approach was used. This was in parallel with continuous testing and upgrading of the CMTS and CM software with bug fixes.

Phase 1 – Activating the OFDMA channel, but not using it for customer data

In this phase we were able to check CMTS and CM behavior, regarding registration, CM management stability, etc. No service flows were assigned to the OFDMA channel. This was done to minimize the customer impact as the upstream traffic was not allowed to use OFDMA channel and instead the traffic continued to use the existing SC-QAM channels. To accomplish this, different techniques had to be used on the two CMTS models deployed in the network. We were able to build the data collectors and



processors, and start acquiring the KPIs selected and RxMER data files, which gives us visibility over the new upstream spectrum and the new OFDMA technology

Phase 2 – Using the OFDMA channel for customer data, but not upgrading their speed tier.

Once the data gathering and modem registration issues were crossed we moved into the phase of using the OFDMA channel for traffic. This phase allowed us to see the real impact of having live data traffic on the channel. Profile/IUC downgrades would happen in one CMTS just based on codeword (FEC) errors. So, with data traffic we could see the quality of the IUCs that we had created and the degree of readiness of the plant. During this phase IUCs were refined and there was a massive amount of work to go out and correct the plant when possible. At this point the customer was not given any speed upgrades, so any channel impairments on the OFDMA meant that the customers would fall back to DOCSIS3.0 capacity and still get the same capacity as they did before.

Phase 3 – Upgrading the upstream speed to customers

Once the IUC definitions were mature and robust, in this phase, the Maximum Sustained Traffic Rate was upgraded in the CM Config files for the eligible customer base. In this phase the impact could be significant as any channel impairments would directly mean less capacity and the customer might not be able to reach the top speed offered in the service.

### **3.3. Technology maturity**

At the time testing and field trialing started, both CMs and CMTSs Upstream implementations were quite immature. There were lots of feature limitations, and some features were not implemented at all, which constrained the roll-out across the footprint. Some examples in the CMTS were a very limited number of IUCs supported, CMTS management of IUCs not reacting to codeword errors, no Upstream PNM implementation, a very limited set of configuration parameters like pilot patterns, subcarrier spacing, cyclic prefix, etc. CMs were more comparatively mature from the beginning as regard support for the OFDMA set of features. Now both CMTS and CMs had numerous bugs that needed to be isolated and solved. This was by far the longest part of the work needed to get OFDMA out in to the field and make it generally available for customers.

Some challenges included issues like handling DBC messages to change IUC definitions as appropriate. Other CMTS devices had issues collecting US RxMER data reliably, and this was a key requirement to understand the plant and create appropriate profiles. Other initial limitations forced Service flows to be assigned to OFDMA channels, even when the operator was not planning on sending data traffic on those channels. Many CMTSs only support 2 IUCs one of which included IUC 13. The profile/IUC definition also only allowed 4 exception zones for modulation order changes, and allowed only one exclusion zone. What this meant was that the allowed configuration options were somewhat limited. Load balancing of US traffic across SC-QAMs and OFDMA channels were also not fully mature.

### **3.4. Initial IUC definition for OFDMA channel**

The CMTS OFDMA channels need some initial IUC definitions to get the upstream channel up and running. For this initial phase of trials and the initial IUC definitions there was no US RxMER information available, as no OFDMA channel was provisioned in the network. Hence a PMA could not be used to create optimized profiles. So, another method was devised to define the initial IUCs for the 6000+ US SGs. The basic idea was to measure the noise on the plant when no SC-QAM or OFDMA channel was present and then use that to estimate the US MER.

### 3.4.1. Data collection for initial IUC

The Upstream data was captured with our data collection system (Viavi XPERTrak), which was continuously collecting signal power measurements in the spectrum where the OFDMA channel was to be located. The measurements were done before the OFDMA channel was turned on. This consisted of an average dBmV level value for each 250kHz bin of spectrum from 5 to 65MHz, collected every 15 minutes. Many days of historical data was used as input to a custom-made algorithm that returned IUC definitions for each US SG.




Figure 6 - Example power level data used to calculate initial IUC definitions

### 3.4.2. Initial IUC Definition

The MER was calculated from the noise level as measured in the signal power measurements in the previous section and the assumption that the CM signal arrives at the CMTS at 0dBmV. For IUC 13 we used the maximum power level found in the data set for each bin as the assumed noise level. For IUC12 we used the average level, choosing a bit more relaxed value for use in designing a higher profile. The power level data had to be tailored and adapted to minislot boundaries, and CMTS specific implementation rules needed to be accommodated on the design of profiles. This algorithm was implemented in a software program that could automatically read the data and output the needed CMTS commands for every serving group.



Figure 7 - Example IUC definition from the deduced MER

The OFDMA features available at the time along with the limitations of each CMTS vendor had to be taken into account for this algorithm to convert this power level measurements into MER and ultimately into modulation orders. Also, worth noting that this method only works before the OFDMA channel is turned on in the plant. As soon as CMs start transmitting in the OFDMA channel, this method cannot be used any longer as the external signal power measurements cannot separate the signal from noise.

### **3.5. Upstream Plant readiness**

Level alarms were setup, in our data monitoring system, for the OFDMA band before OFDMA signals were activated. Whenever the threshold conditions were crossed, technicians were dispatched to check the plant, and correct noise sources or pinpoint customer installation problems. This happened for about 8% of the US SGs. Each and every US SG was again verified and acted upon just before provisioning a customer with a higher top speed that actually needs the OFDMA channel capacity.

Based on the KPIs described below, a technician truck roll was ordered for the SGs with high number of impairment-hours. Technicians would go through the plant, amplifier by amplifier, tap by tap, looking for the source of the noise impairment and trying to isolate it. If the source of the problem happened to be in the outside plant, the technicians would take steps to solve it. If the source was identified to be inside a home, a filter blocking part of or the complete upstream would be used on the drop, and then a home network maintenance work item would be scheduled.

This is very similar to the normal maintenance of the upstream in DOCSIS3.0. The main difference here was the large amount of work it took to do this for the whole network in a short amount of time. The other difference from DOCSIS 3.0 is also the use of this new spectrum which is in the lower part of the upstream spectrum and very noisy. As with the legacy 3.0 spectrum, this is a process which will continue on a daily basis, it is now even more demanding as this new lower portion of the spectrum is more susceptible to noise and interference. This will continue to be ongoing operations work, just like it was for the upper part of the spectrum before OFDMA. High pass filters that used to be deployed to block noise can no longer be used in the OFDMA region, making it harder and more expensive to maintain the outside plant.

### **3.6. KPIs for monitoring**


We defined a new KPI to monitor OFDMA upstream channels: IUC-usage-hours. IUC-usage-hours is defined by number of hours in a day a CM uses each IUC on a particular OFDMA channel. This quickly became the main KPI that we tracked during these initial deployments and Upstream OFDMA roll-out. We tracked this metric of IUC-usage-hours over days, weeks and now months. As an example, a CM could be on IUC 12 for 22 hours a day and IUC 13 for 2 hours a day, while another CM perhaps could be in an OFDM-impaired state (partial service) for an hour a day.

Now in order to understand the reason for the variance in IUC-usage-hours, a homemade tool was built with the following additional KPIs:

- Codewords per IUC (total, errored): Per CM statistics on codeword errors, to understand the customer impact.
- Receive and Transmit Power: Per CM Rx/Tx power
- Data volume / Heavy user: The data volume metric was considered useful to diagnose different CM behaviors in the same US SG, as CMTS depends on codeword error ratios (CER) for IUC


management and its algorithm depends on traffic. Heavier users are more likely to hit thresholds and experience IUC changes.

These set of metrics proved essential, as this was the data used to fine tune IUCs, to make decisions along the trials and subsequent rollouts, to order plant cleaning or to dispatch a technician to the customer premises.



**Figure 8 – Tracking IUC-usage-hours per day (across CMTS)**

CMTSs automatically change the IUC used for each CM based on CER thresholds. The effect of one of the adjustments made during the process of deploying OFMDA can be seen in Figure 8 above. When changing the IUC retry interval from 180 minutes to 90 minutes, the percentage of time CMs spent in IUC 12 and 13 increased, and the time spent without OFDMA capacity decreased. At this point in time (i.e. prior to the PMA field trial) the CMTS IUC change mechanism algorithm was flawed and customers could sometimes have packet loss for a few seconds during IUC downgrade. So, the goal became not wanting to have many IUC switches. This testing allowed the operations team to decrease this time without any noticeable service degradation.



14 – denotes Impaired OFDMA operation

| Cordas com maior % em IUC13 + IUC14 - UPSTREAM |        |             |             |           |            |        |        |        |              |
|------------------------------------------------|--------|-------------|-------------|-----------|------------|--------|--------|--------|--------------|
| Ano                                            | Semana | Corda       | CMTS        | Total CMs | # Amostras | %IUC12 | %IUC13 | %IUC14 | %IUC13 ou 14 |
| 2020                                           | 33     | VCH13 C27D2 | PNO1CMTS013 | 12        | 1884       | 34     | 17     | 49     | 66           |
| 2020                                           | 33     | PNF03 3JD1  | PDS1CMTS004 | 9         | 1502       | 22     | 29     | 48     | 78           |
| 2020                                           | 33     | PTM25 2JD1  | PTM1CMTS005 | 24        | 3666       | 35     | 18     | 47     | 65           |
| 2020                                           | 33     | AGU03 C41D2 | QUETCMTS012 | 15        | 2518       | 38     | 22     | 40     | 62           |
| 2020                                           | 33     | SES01 C35D1 | FGT1CMTS005 | 4         | 672        | 43     | 19     | 38     | 57           |
| 2020                                           | 33     | PVZ05 3JD1  | VCD1CMTS003 | 11        | 1815       | 40     | 25     | 35     | 60           |
| 2020                                           | 33     | TMA10 2JD1  | FGT1CMTS004 | 20        | 3144       | 23     | 42     | 35     | 77           |
| 2020                                           | 33     | BBA05 C35D1 | PNO1CMTS009 | 6         | 1004       | 31     | 35     | 33     | 69           |
| 2020                                           | 33     | SJM02 C53D1 | RME1CMTS003 | 2         | 504        | 67     | 0      | 33     | 33           |
| 2020                                           | 33     | FIG01 3JD2  | FIG1CMTS005 | 11        | 1847       | 42     | 25     | 33     | 58           |
| 2020                                           | 33     | PAR04 3JD1  | PAR1CMTS006 | 23        | 3751       | 67     | 0      | 33     | 33           |

**Figure 9 – Operational dashboards displaying CM IUC distribution in one US SG and top US SGs with CM impairments**

### **3.7. US RxMER Data Collection from the CMTS**

The DOCSIS 3.1 CCAP OSSI specification defines a standard procedure to use SNMP to start upstream RxMER measurement on the CMTS and have the encoded binary files that comply to the specification defined format uploaded to a TFTP server. However, this feature is not currently supported by all of the available CMTSs. In our field data collection practice, 2 other methods were used to collect upstream RxMER data based on what the CMTSs support.

#### **3.7.1. Standard SNMP + TFTP method**

The DOCSIS 3.1 CCAP OSSI specification defines MIBs and data formats to allow upstream RxMER data collection. The high-level steps are described below, on a per CM basis

- The data collector specifies the CM MAC address and CM upstream channel interface number, sets parameters for upstream RxMER data measurement, and specifies the TFTP server address and directory for file upload through the MIBs on the CMTS
- The CMTS performs upstream RxMER measurement for the specified CM and OFDMA channel
- The CMTS uploads the encoded file that contains the upstream OFDMA RxMER measurement data to the specified TFTP address and directory
- The data collector reads the uploaded binary file, decodes it and stores it into the data service for PMA profile calculation

#### **3.7.2. SNMP + SFTP method**

Some CMTSs may comply to part of the DOCSIS 3.1 CCAP OSSI specification and may support the SNMP MIBs for upstream RxMER measurement. However, they may require different methods to gather the collected data, such as SFTP. The steps are described below, on a per CM basis

- The data collector specifies the CM MAC address and CM upstream channel interface number, sets parameters for upstream RxMER data measurement through the MIBs on the CMTS
- The CMTS performs upstream RxMER measurement for the specified CM and OFDMA channel
- The CMTS stores the encoded file that contains the upstream OFDMA RxMER measurement data in the CMTS's local storage (the CMTS may only store a limited number of PNM files)
- The data collector retrieves the stored binary file through SFTP from the CMTS, decodes it and stores it into the data service for PMA profile calculation

#### **3.7.3. CLI method**

Some CMTSs may not support the MIBs for upstream data collection that are defined in the DOCSIS 3.1 CCAP OSSI specification. However, they may support measuring the upstream OFDMA RxMER data and presenting the data through the command line interface (CLI). In this case, we can parse the CLI output to collect the upstream OFDMA RxMER data. The steps are described below, on a per CM basis

- The data collector uses a CLI client to specify the CM MAC address and CM upstream controller/interface/channel number in the CMTS CLI and collects the CLI output (the CMTS may not support data collection triggering, instead, it may have a time interval to be configured to periodically perform upstream RxMER measurement)
- The CLI client automatically parses the CLI output and generates JSON formatted upstream OFDMA RxMER data
- The data collector stores the JSON formatted data into the data service for PMA profile calculation

### **3.8. US RxMER Data Collection in field deployment for PMA trials**

CableLabs DOCSIS Common Collection framework (DCCF) software can collect PNM data from the plant. However, the CMTSs were not fully compliant to PNM specification for PNM testing and reporting. One CMTS could generate the RxMER files and store them in its filesystem, but it didn't implement the TFTP upload. So, for the sake of this field trial, some shell scripts were developed to trigger RxMER measurement through SNMP and fetch the encoded binary files from the CMTS using SFTP.

Data was collected for all the CMs and all the SGs in field trial every 2 hours, for a couple of weeks. A software limitation and the total number of CMs imposed this 2-hour limit. In total, more than 13,000 RxMER files were collected. For practical reasons, all SGs we collected data belonged to the same CMTS.

## **4. US PMA Algorithms**

The upstream RxMER data has different features compared to what we see in the downstream. Although one could start with the downstream PMA's clustering algorithm to calculate upstream OFDMA modulation profiles, given the difference of the Upstream RxMER data (as explained in section 2.9) it become clear that a simple "clustering of CMs" approach would not be a good solution. Due to noise funneling since the US RxMER were similar many of the profiles are expected to have similar characteristics. So, we developed several new candidate algorithms for comparison and improving the robustness of the upstream. The immediate goal for this initial rollout was focused around developing IUCs which modems could stay on without profile flapping. Each profile is expected to be used by a group of CMs that have similar channel characteristics.

Each of the Minislots in the Upstream OFDMA channel can be configured to use a different modulation order. This allows the operator to optimize the upstream transmissions across the wide frequency band (10-96 MHz) of the channel. The specific choice of modulation order selected for each minislot is communicated to the CMs in the form of an IUC (modulation profile) which allows the CM to modulate the signal accordingly. An IUC/modulation profile consists of a vector of bit-loading values, an integer value for each active Minislot in the upstream channel. Since the modulation orders range from QPSK to 4K-QAM, the range of bit-loading values is from 2 to 12.

The PMA generates an IUC 13 that is the lowest common denominator profile, which can be successfully used by all CMs in the Service Group. A CMTS can support up to 7 modulation profiles, including IUC 13. Each CM can be assigned up to 2 modulation profiles at a time, including IUC 13 and an optimized profile for the CM's unicast traffic.

This capability, the ability to optimize the upstream transmission for the channel characteristics of the CM population, is a powerful feature that allows for a significant improvement in robustness and channel capacity. The CMTS and CM perform measurements and report network conditions as a part of supporting PNM functionality in the DOCSIS network. The DOCSIS 3.1 Upstream PNM Measurements includes: US active and quiet probes, Triggered spectrum capture, US equalizer coefficients, Impulse noise statistics FEC statistics, Histogram, Channel Power and the RxMER per subcarrier. So far, we are basing the PMA profile creation algorithms on the US RxMER data, in the future one can include other upstream data sets to fine tune the profiles that we create.

## 4.1. Differences from Downstream Algorithms

The downstream PMA algorithm looks into the variances among the CMs, clusters the CMs into groups, and assign different modulation profiles to each CM group. The algorithm can use a data snapshot that's captured only once from the CMs, or it can use pre-processed data (average, minimum, percentile etc.) from each CM's multiple historical data captures.

On the upstream, each CM's RxMER data tend to have similar patterns when they are captured during the same time slot because of noise funneling. The variance is much more within the CM's RxMER captures over a relatively long period of time. Hence, the upstream PMA algorithm focuses more on optimizing the channels robust operation and CMs' upstream capacity over a certain period (when the operator does not plan to change upstream profiles frequently such as changing the profiles every few hours). The candidate upstream PMA algorithms also consider the fact that the CMTSs automatically upgrade/downgrade the upstream modulation profiles for each CM based on their monitored FEC performance, which makes the time clustering based profiles more effective.

Problem: Given a set of US RxMER data samples from CMs, return optimized profile definitions.

We have two classes of algorithms we developed and are field trialing, one is developing using the percentile method and the other is using time clustering methods. These methods and their variants are described below.

## 4.2. Percentile method

The CMTS is actively managing the CM's IUC/profile assignment as plant conditions change. The upgrades/downgrades to the CMs' upstream OFDMA modulation profile is based on the CMs' FEC performance or RxMER data. The percentile method for creating profiles is a simple statistical set of methods which can create robust profiles based on the past performance of the plant. The idea is to choose a conservative profile which can fit most of the CMs, most of the time. So, the algorithms arrange the CMs in descending order of their RxMER values and choose to design a few different profiles at a few different percentile values (e.g. 0.5 percentile and 2 percentile)

This method focusses on optimizing the overall channel robustness so that an CMTS can maintain good upstream service for most of the CMs.

### 4.2.1. Algorithm 1A : Per CM Percentile

Inputs: A list of CM RxMER per subcarrier, choice of a percentile numbers that an operator wants to choose at the per CM level and for the profile level.

Outputs: List of robust profile definitions for use on the upstream channel

Algorithm:

- Calculate a representative CM RxMER sample for each CM from the data captured over time
  - For each CM on the US channel
    - Create an artificial 'x<sup>th</sup>' percentile sample of RxMER for this CM
      - Start with the first sub carrier in the channel
      - Find the 'x<sup>th</sup>' percentile RxMER value across all the samples for that sub carrier for that CM

- Repeat for all subcarriers
  - Repeat for each ' $x^{\text{th}}$ ' percentile chosen by the operator.
    - E.g. 0.5<sup>th</sup> percentile and 2<sup>nd</sup> percentile could be the two percentile values for ' $x^{\text{th}}$ '. The number of percentile values depends on the number of needed IUCs.
- Create profiles from CMs' representative RxMER samples (percentile samples)
  - Group each of the ' $x^{\text{th}}$ ' percentile data samples from the CMs (from the previous step)
  - Create new ' $y^{\text{th}}$ ' percentile (normally 0.5% or 0%) samples across each of the CMs' ' $x^{\text{th}}$ ' percentile groups.
    - Start with the first subcarrier in the channel
    - Find the ' $y^{\text{th}}$ ' percentile RxMER value across all the CM samples from an ' $x^{\text{th}}$ ' percentile group
    - Repeat for all subcarriers
    - Repeat for all ' $x^{\text{th}}$ ' percentile groups
  - Translate each of the ' $y^{\text{th}}$ ' percentile RxMER to Modulation orders per [PHYv3.1] spec.

#### **4.2.2. Algorithm 1B: All CMs Percentile**

Inputs: A list of CM RxMER per subcarrier, choice of a percentile numbers that an operator wants to use to create profiles.

Outputs: List of robust profile definitions for use on the upstream channel

Algorithm:

- Create a profile from all of the RxMER samples from all of the data
  - Create a new ' $y^{\text{th}}$ ' percentile sample across all of the CM's RxMER samples.
    - Start with the first sub carrier in the channel
    - Find the ' $y^{\text{th}}$ ' percentile RxMER value across all the samples for all the CMs
    - Repeat for all subcarriers
  - Repeat for as many ' $y^{\text{th}}$ ' percentile values as many as IUCs are needed.
    - E.g. IUC 13 could be the 1<sup>st</sup> percentile, and IUC 12 could be the 5<sup>th</sup> percentile.
  - Translate each of the ' $y^{\text{th}}$ ' percentile RxMER to Modulation orders per the PHY spec.

#### **4.2.3. Algorithm 1C: Remove Outliers + Percentile**

Inputs: A list of CM RxMER per subcarrier, choice of a percentile numbers that an operator wants to use to create profiles, and a choice of how to detect outlier CMs.

Outputs: A list of robust profile definitions for use on the upstream channel

Algorithm:

- Find and remove Outlier CMs.
  - Method C1
    - For each CM, across all of its samples: Calculate an average RxMER number for the CM
    - Calculate Standard deviation for each CM, with respect to all of its samples.
    - Remove CMs outside the Confidence interval (e.g.: 98%) lower bound.
      - Remove all samples of all outlier CMs from the data set



- Method C2
  - Calculate average RxMER: Each CM, all samples
  - Remove CMs below 70% of average value (of all CMs)
    - Remove all samples of these CMs ( in some cases, this may be an empty operation), from the data set.
- Create a profile set from the remaining RxMER samples
  - Use Algorithm 1A or 1B

### 4.3. Time Clustering Methods

The CMTS automatically upgrades/downgrades the CMs' upstream OFDMA modulation profile based on the CMs' FEC performance. The hypothesis for the time clustering method is that we can find clusters of CMs with similar RxMER over time and design IUCs based on those clusters to keep the CMs on IUCs for long periods of time. The CM clustering algorithm that's been previously developed and used in downstream PMA is repurposed to calculate clusters from the data captured over time on all CMs. The time clustering methods focus on optimizing the overall channel capacity through days or weeks of time without recalculating or modifying the upstream modulation profiles very frequently. The J value that's used in the downstream for capacity gain measurement is repurposed to measure the theoretical capacity gain from the time clustering based profiles over a time period.

#### 4.3.1. Algorithm 2A: Time Clustering using Actual CM Samples

Inputs: A list of CM RxMER per subcarrier, choice of a number of profiles an operator.

Outputs: A list of robust profile definitions for use on the upstream channel

Algorithm :

- Find Time clusters from all CM's and all of their samples
  - Use actual sample of RxMER for every CM, every measurement
  - Use a clustering algorithm to find groupings across full-RxMER values.
    - e.g. reuse the PMA Algorithm (for Downstream)
    - Find 2-7 clusters
- For each cluster/group of CM-samples
  - Choose an RxMER value per subcarrier, from this groups samples as follows
    - The average value at each subcarrier
    - The minimum value at each subcarrier
    - A certain percentile value (picked by the operator)
    - Some other centroid definition
  - Translate each of the resulting RxMER vector to Modulation orders per the PHY spec.

#### 4.3.2. Algorithm 2B: Time clustering using artificial Percentile Samples

Inputs: A list of CM RxMER per subcarrier, choice of a number of profiles an operator.

Outputs: A list of robust profile definitions for use on the upstream channel

Algorithm:

- For All CM's and all samples

- Create 199 artificial samples of RxMER for each CM as described below
  - An RxMER value per subcarrier is chosen using a percentile value across all samples of that CM
  - Samples will use percentiles from 0.5% to 99.5% in increments of 0.5%
- Perform clustering with these artificial samples
  - Use Algorithm 2A (This step will find the most common percentiles)

#### **4.3.3. Algorithm 2C: Time Clustering after removing outliers**

- For All CM's and all samples
  - Use all samples of RxMER for each CM
  - Create an artificial threshold sample using a percentile value (e.g. 1%)
    - RxMER value per Subcarrier is chosen using a %tile value across all samples
  - For every sample
    - If any individual RxMER within the sample is below artificial threshold sample, then bump up those RxMER values to the threshold MER level using a ceiling-like calculation. This keeps the weight of the sample for other areas in the channel which are unchanged
    - Instead of bumping up samples, we could also ignore the samples
- Now with the outliers removed, use Algorithm 2A.

#### **4.4. Other PMA Considerations**

Modulation order of profile from minislot RxMER: There are a few different methods to choose the Modulation order of the profile from the RxMER values of all the subcarriers within that minislot. One could choose one of the following RxMER values to use when translating to the modulation orders

- Average RxMER of subcarriers within minislot.
- Majority RxMER of subcarriers within minislot
- X percentile RxMER of subcarriers within minislot
- Minimum RxMER of subcarriers within minislot

Margins when translating from RxMER to modulation order: Different margins can be applied at different frequencies within the channel. Lower frequencies could have a higher margin, while higher frequencies wouldn't need as much of a margin. Another way to apply different margins is to preprocess the US RxMER data and if there is a higher variation in the data, one can apply a higher margin. This way the operator can lower the US RxMER before sending to PMA algorithm

Capacity gain Calculation/Optimization function: The downstream has a J-Value calculation which can calculate the capacity achieved by a set of profiles. In the downstream the profiles are also weighted by the number of CMs associated with the profile to calculate capacity. Also, the gain is calculated with reference to 256-QAM in the downstream. For the upstream given the dynamic nature of the noise on the plant and the fact that many of the CMs have the same RxMER signature, the idea of an optimization function can be built around how much time is spent on each of the IUC/profiles, i.e. of the many RxMER data samples of all CMs fit which a particular profile, we can weight each profile by that percentage of samples, to calculate an overall weighted capacity. To figure out if an US RxMER sample can fit a profile, the idea would be to establish a threshold across all subcarriers: e.g. only 1 or 2 % of subcarriers can be below the required RxMER level for that modulation order. Also, this comparison should be done on a minislot basis. A good reference in the upstream to capture the capacity gain would be to compare to 16-QAM or 64-QAM, depending on the plant.

PMA Exception handling: Due to adjacent channel interferences (from adjacent SC-QAMs), there is a need to handle exceptions for the Higher and lower edges of the OFDMA channel. The idea is that a few number of minislots (2-4) need to be limited to one or two lower modulation orders than normal to account for interference from an adjacent SC-QAM

## 5. US PMA Field Trial Results

### 5.1. US RxMER Field Data

The below figure 10 shows a sample of the US RxMER collected at the CMTS, for 4 cable modems within one Serving group. The data was collected for every cable modem, every 2 hours for over 3 weeks.

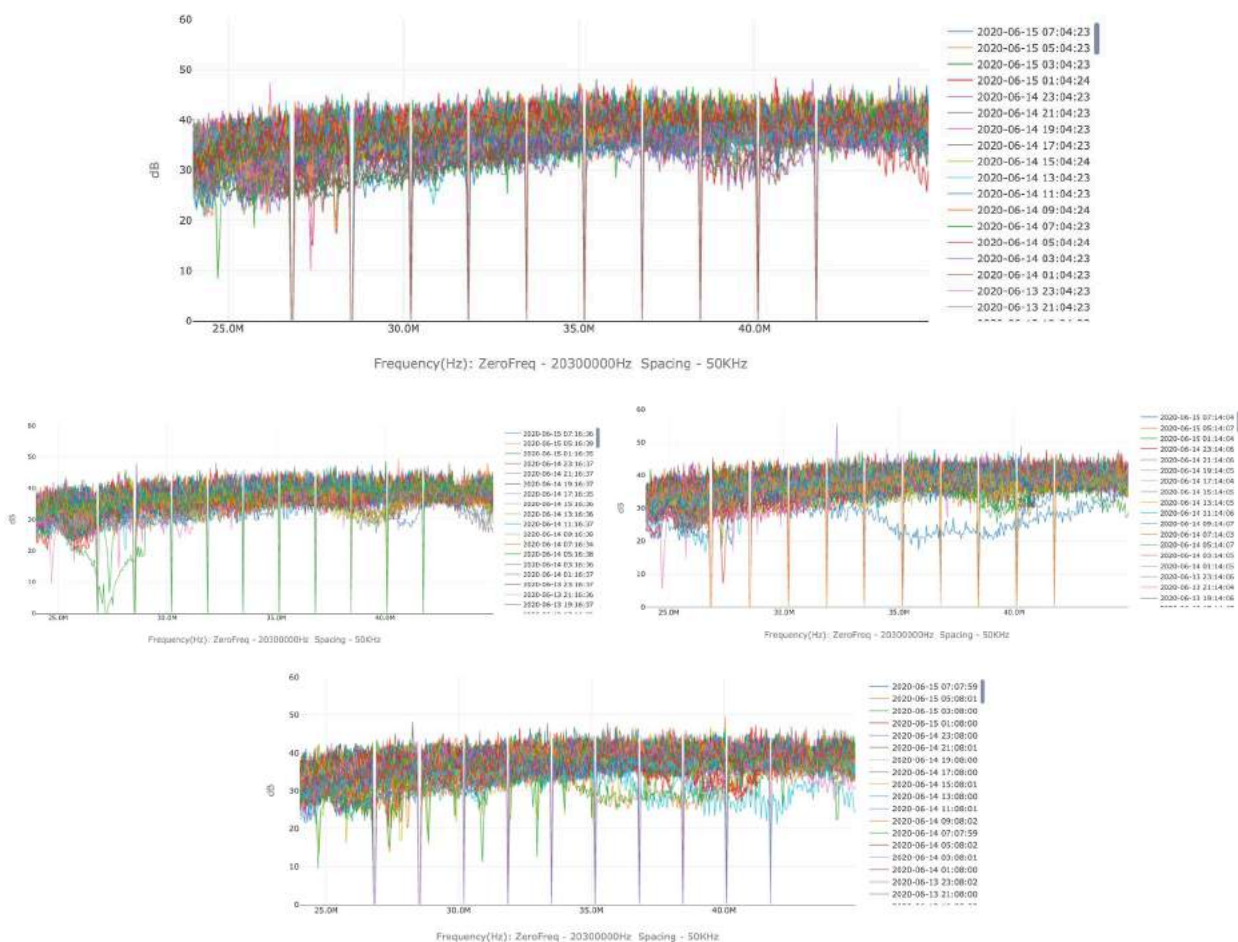
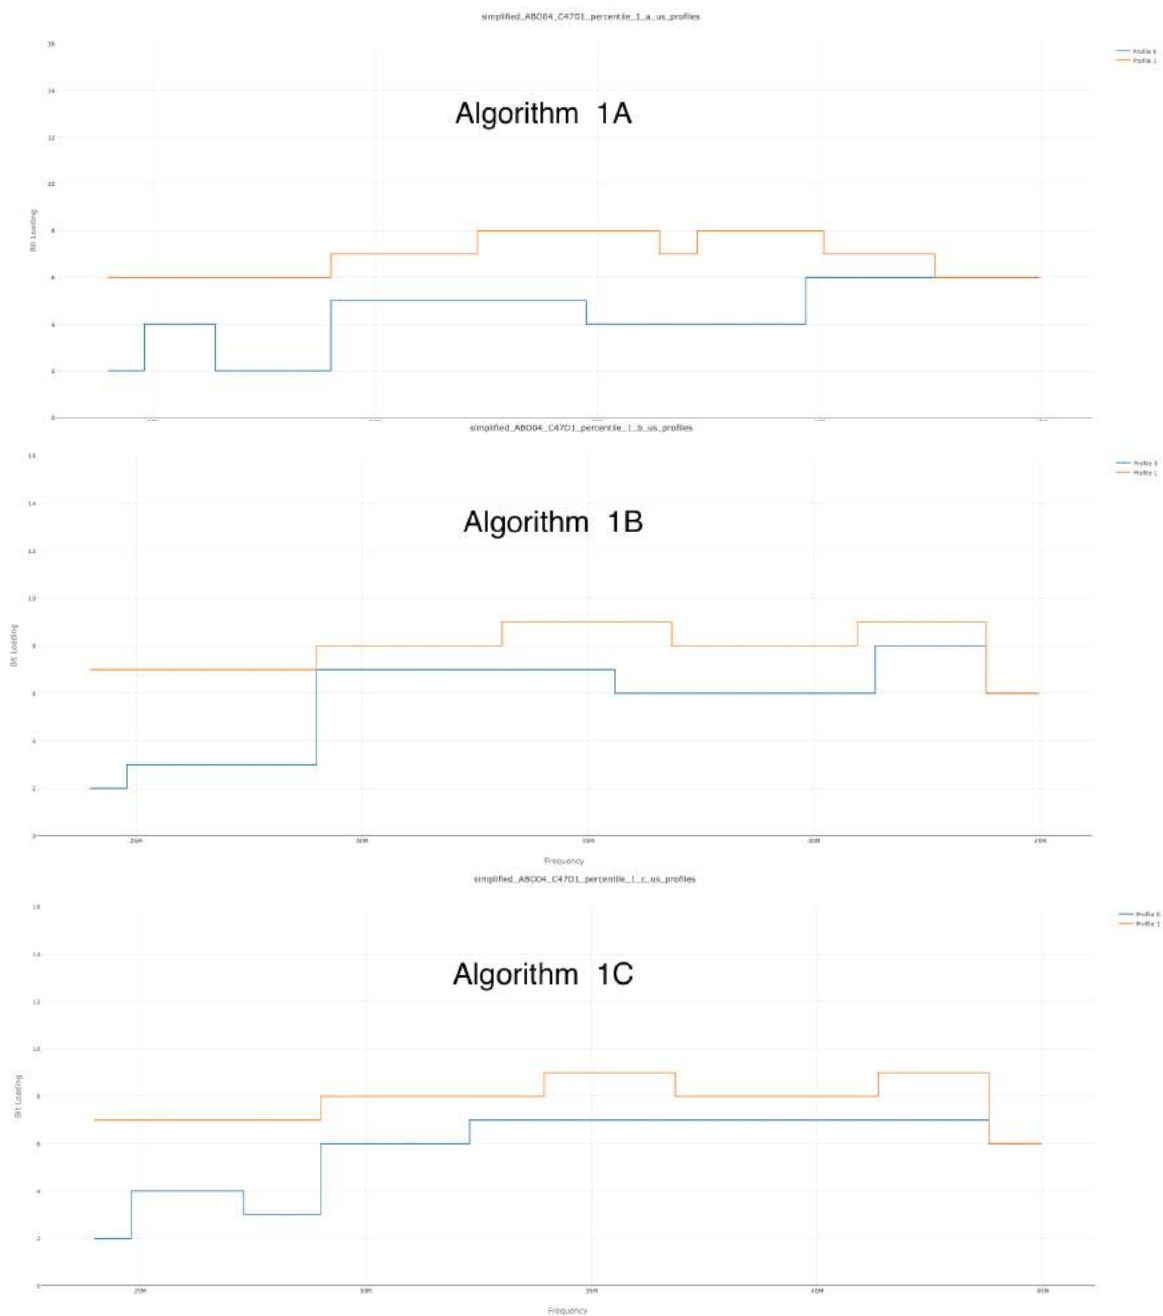


Figure 10 – Us RxMER Samples from 4 CMs


### 5.2. US IUCs/Profiles designed

The following figures show the profiles designed by the Algorithm 1A,1B,1C, 2A, 2B,2C from the section above. This iteration of profile creation was limited to 2 IUCs due to CMTS limitations. The profile 0 in blue is used as IUC 13 and the profile 1 in orange is used as IUC 12.



**Figure 11 – IUCs created using Percentile Algorithms 1A,1B,1C**

The profiles created by algorithm 1A are more conservative than the 1B and 1C as those algorithms remove the outliers or ignore the worst samples across all the CMs equally.



**Figure 12 – IUCs created using Time Clustering Algorithms 2A, 2B, 2C**

The profiles created by algorithm 2A is more conservative than the 2B as 2A operates on all the original data samples, whereas algorithm 2B operates on a subsampled percentile data set. Algorithm 2C is somewhat in between as it removes the outliers but keeps the weighting of all the samples.

The goal for the field trial is to try out each of these 6 methods of profile creation and uses the profiles they create in the field for a week and observe the performance of these profiles.

### 5.3. Performance and Stability

The table below depicts the field results for the profiles designed using algorithms 1A and 1B.

**Table 4 – Field Results from using PMA Algorithm 1A, 1B**

|                                       | capacity (Mbps/MHz) |       |              |       | Performance  |       |          |              |       |          |
|---------------------------------------|---------------------|-------|--------------|-------|--------------|-------|----------|--------------|-------|----------|
| <i>Upstream<br/>Serving<br/>group</i> | Algorithm 1A        |       | Algorithm 1B |       | Algorithm 1A |       |          | Algorithm 1B |       |          |
|                                       | IUC12               | IUC13 | IUC12        | IUC13 | IUC12        | IUC13 | impaired | IUC12        | IUC13 | impaired |
| ABO04<br>C47D1                        | 5.9                 | 3.7   | 6.8          | 4.8   | 97.8%        | 1.6%  | 0.6%     | 96.4%        | 3.4%  | 0.2%     |
| PAR06<br>2JD1                         | 4.6                 | 2.4   | 4.1          | 3.1   | 99.0%        | 0.8%  | 0.2%     | 95.9%        | 3.8%  | 0.3%     |
| PSL03<br>C53D2                        | 4.4                 | 2.3   | 4.8          | 3.6   | 92.0%        | 3.6%  | 4.4%     | 93.8%        | 4.5%  | 1.7%     |

The Table above shows the IUC-usage-hours percentage for the period and the channel capacity in Mbps per MHz for each algorithm.

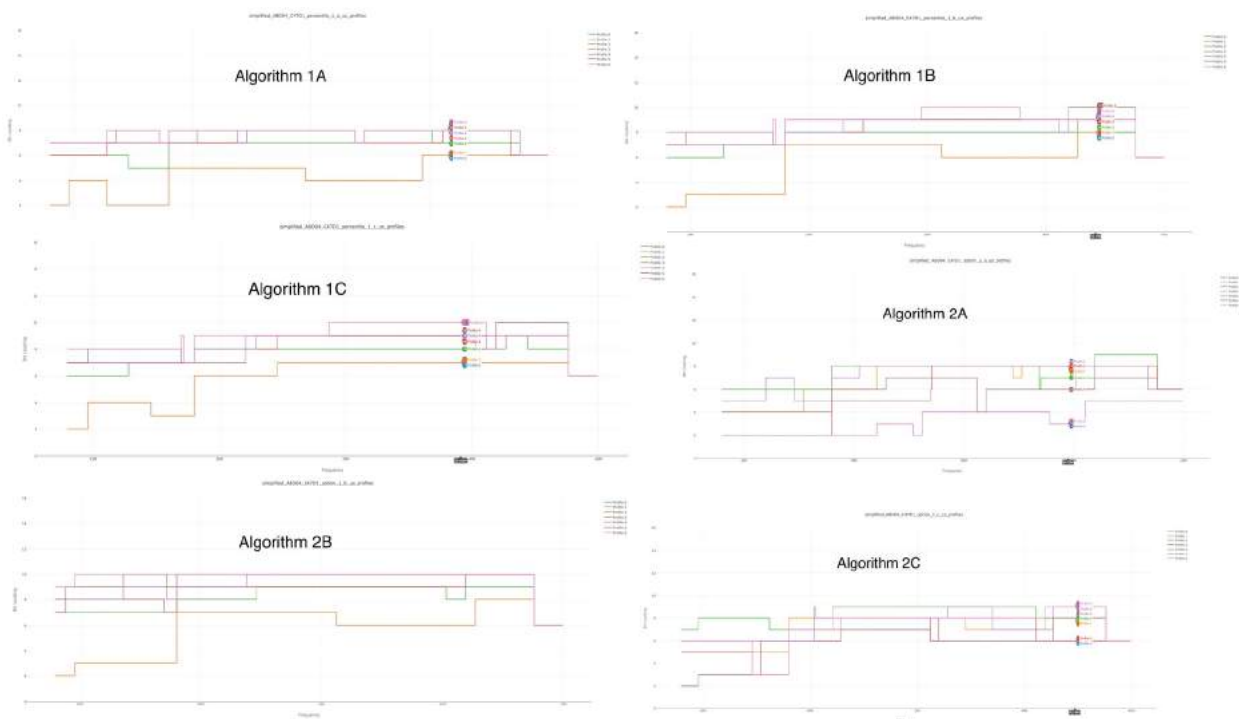
IUC-usage data was collected each hour. This CMTS upgraded IUCs without any other consideration X minutes after it was downgraded and X was configured in this trial to 15 minutes. This means the results in the table are a sampling of the actual behavior, but still a good picture of the reality as each set of profiles (from each of the algorithms) was run for more than one week.

CMTS manages IUCs based on FEC errors. When the number of errored codewords crosses a threshold on a configurable codeword window, the IUC is downgraded (from 12 to 13, and from 13 to no-IUC, i.e., channel impaired). So, at this point we are not tracking FEC error rates because the CMTS will essentially downgrade the IUC up to the point that there are no errors.

Due to the working restrictions due to Covid, the field trials were delayed and slowed down a bit, and at the time of writing, we have not had a chance to evaluate the other 4 methods of profile creation. So far Algorithm 1A shows greater stability with the designed IUC 12. Our hypothesis is that profiles based on methods 2A,2B,2C will be likely give us more stable upstream operations. We hope to share those results in the near future.

### 5.4. US IUCs/Profiles designed with no IUC limitations

We ran the same algorithms on the data from the serving groups to create more than 2 IUCs, we chose 7 IUCs as that will be maximum number supported by the CMTS per channel, in the future. The figure 13 below gives us an idea of the profiles created for a single serving group. In these Serving groups we currently have only ~30 CMs each, so some of the profile definitions are overlapping (we chose percentile values of 0.5,3,6,9,12,15,18.) For a larger set of CMs we expect the profile definitions to be more spread out.



**Figure 13 – 7 IUCs created using all Algorithms**

## 6. Conclusion / Future Work

OFDMA is a newer technology with a lot of promise for operators looking to increase their upstream bandwidth capacity. At this point, OFDMA is a bit of a long walk to the finish line, because of novelty, the immature implementations of the technology (which is quickly being fixed) and the use of new-noisy spectrum which hasn't been used before. There will be a cyclic process of cleaning up the upstream plant, followed by degradation and new issues seen in the plant followed by more plant work and upgrades.

As it stands today, more flexibility in the CMTS OFDMA system functionality is needed, as this will enhance the upstream stability and ultimately the customer experience. There are many CMTS limitations, the most important of them being the limitation around just 2 IUCs per channel. The number of IUCs supported per channel needs to increase to accommodate the upstream variations. Improvements in the ability to define more than the limited (currently 4) modulation orders per profile would help design more granular profiles and improve the stability of the modems. There is an immediate need to measure and use the US RxMER to design and then select IUCs, and the use of FEC errors to downgrade IUCs. A fuller implementation of Upstream PNM functionality as defined in the DOCSIS 3.1 specifications will go a long way to ensuring a smooth transition from SC-QAM upstreams to OFDMA upstreams.

Lower frequencies in the upstream can be a very noisy and can make the channels using that portion of the spectrum unstable. DOCSIS 3.1 brought a high degree of adaptability with the new OFDMA technology. CMTSs can switch between upstream profiles (IUCs) to cope with the variability of the plant but the operator needs to create and design the IUCs optimally. So, a Profile management Application (PMA) is even more necessary on the upstream. The work done here shows some encouraging results for different upstream PMA algorithms and further testing is ongoing.

In terms of future work, we plan to complete the analysis on all the profile creation techniques and share the top methods with industry. Also, when more IUCs per channel are available, it would be interesting to see if which methods can be used to create better profiles which can increase capacity and keep the upstream operation robust. We also plan to experiment with pilot patterns and other OFDMA parameters like the number of symbols per frame, FFT size, Cyclic prefix, choosing exclusion bands etc. and see how best to optimize the OFDMA channel operation.

## Abbreviations

|       |                                               |
|-------|-----------------------------------------------|
| CW    | Codeword                                      |
| CMTS  | Cable modem termination system                |
| CM    | Cable modem                                   |
| CER   | Codeword error rate                           |
| DS    | Downstream                                    |
| DCCF  | DOCSIS Common Collection Framework            |
| HFC   | hybrid fiber-coax                             |
| Hz    | hertz                                         |
| FEC   | forward error correction                      |
| FFT   | Fast Fourier transform                        |
| IDFT  | Inverse discrete Fourier transform            |
| IUC   | Interval Usage Code                           |
| KPI   | Key Performance Indicator                     |
| SG    | Service group                                 |
| SID   | Service Identifier in the DOCSIS upstream     |
| OFDM  | Orthogonal Frequency Division Multiplexing    |
| OFDMA | Orthogonal Frequency Division Multiple Access |
| PMA   | Profile Management Application                |
| PNM   | Proactive Network Maintenance                 |
| RxMER | Receive Modulation Error ratio                |
| US    | Upstream                                      |
| SCTE  | Society of Cable Telecommunications Engineers |

## Bibliography & References

*[MULPIv3.1] DOCSIS MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I20-200407, April 7, 2020, Cable Television Laboratories, Inc.*

*[PHYv3.1] DOCSIS Physical Layer Specification, CM-SP-PHYv3.1-I17-190917, September 17, 2019, Cable Television Laboratories, Inc.*

*[CCAPOSSv3.1] DOCSIS Converged Cable Access Platform Operations Support System Interface Specification, CM-SPCCAP-OSSv3.1-I18-200610, June 10, 2020, Cable Television Laboratories, Inc.*

*DOCSIS 3.1 Profile Management Application and Algorithms, 2016 Spring Technical Forum Proceedings, Greg White and Karthik Sundaresan (CableLabs)*

*Practical Lessons from D3.1 Deployments and a Profile Management Application (PMA), 2019 SCTE Expo.*



# **Latency Measurement: What is latency and how do we measure it?**

A Technical Paper prepared for SCTE•ISBE by

**Karthik Sundaresan**  
Distinguished Technologist  
CableLabs  
858 Coal Creek Circle, Louisville, CO, 80303  
3036613895  
k.sundaresan@cablelabs.com

**Greg White**  
Distinguished Technologist  
CableLabs  
858 Coal Creek Circle, Louisville, CO, 80303  
3036613822  
g.white@cablelabs.com

**Steve Glennon**  
Distinguished Technologist  
CableLabs  
858 Coal Creek Circle, Louisville, CO, 80303  
3036613834  
s.glennon@cablelabs.com

# Table of Contents

| Title                                                             | Page Number |
|-------------------------------------------------------------------|-------------|
| 1. Introduction.....                                              | 4           |
| 1.1. Quality of Experience .....                                  | 4           |
| 1.2. Latency in the Internet.....                                 | 5           |
| 1.3. Common techniques to reduce latency.....                     | 5           |
| 2. View of latency measurements .....                             | 6           |
| 2.1. MSO Goals for Latency Measurement.....                       | 6           |
| 2.2. Current National Latency Reports.....                        | 7           |
| 2.2.1. Measuring Broadband America .....                          | 7           |
| 2.2.2. Measuring Broadband Canada Project.....                    | 8           |
| 2.2.3. EU Broadband Report.....                                   | 8           |
| 2.2.4. Speedtest (Ookla) .....                                    | 8           |
| 3. Latency Metrics .....                                          | 9           |
| 3.1. One-way Latency (or Packet Delay) .....                      | 9           |
| 3.2. Round Trip Latency .....                                     | 9           |
| 3.3. Singleton Measurements vs Sets of Measurements .....         | 10          |
| 3.4. Jitter or Delay Variation .....                              | 10          |
| 3.4.1. Inter-Packet Delay Variation, IPDV .....                   | 11          |
| 3.4.2. Packet delay variation, PDV .....                          | 11          |
| 3.4.3. Jitter Metrics in Use in industry .....                    | 11          |
| 3.5. Descriptive statistics.....                                  | 12          |
| 3.5.1. Basic statistics.....                                      | 12          |
| 3.5.2. Percentile Numbers .....                                   | 13          |
| 3.6. Histograms .....                                             | 14          |
| 3.7. Visualization of Latency .....                               | 14          |
| 3.7.1. Time series.....                                           | 14          |
| 3.7.2. Probability density function (PDF).....                    | 15          |
| 3.7.3. Cumulative Distribution Function, CDF .....                | 15          |
| 3.7.4. Complementary cumulative distribution function, CCDF ..... | 15          |
| 3.7.5. Example PDF/CDF/CCDF .....                                 | 16          |
| 4. Latency Measurement architectures .....                        | 18          |
| 4.1. Types of measurement .....                                   | 18          |
| 4.1.1. Active measurements .....                                  | 18          |
| 4.1.2. Passive measurements.....                                  | 18          |
| 4.2. Industry measurement architectures.....                      | 19          |
| 4.2.1. SamKnows Whitebox (dedicated test device solution) .....   | 19          |
| 4.2.2. The M-Lab NDT (User initiated).....                        | 20          |
| 4.2.3. TWAMP .....                                                | 20          |
| 4.2.4. Simple Two-Way Active Measurement Protocol.....            | 21          |
| 4.3. Measurement considerations .....                             | 22          |
| 4.3.1. Measurement under load vs quiet times .....                | 22          |
| 4.3.2. Window over which the measurement is done. ....            | 22          |
| 4.3.3. Off-net and On-net testing .....                           | 23          |
| 4.3.4. Latency Measurement Test definitions .....                 | 23          |
| 4.3.5. Marked traffic vs Unmarked traffic. ....                   | 23          |
| 5. Conclusion.....                                                | 24          |
| Abbreviations .....                                               | 24          |
| Bibliography & References.....                                    | 25          |

## List of Figures

| <b>Title</b>                                                                           | <b>Page Number</b> |
|----------------------------------------------------------------------------------------|--------------------|
| Figure 1 – MSO view of Latency Measurements .....                                      | 6                  |
| Figure 2 – One way Latency vs. Round trip Latency .....                                | 9                  |
| Figure 3 – Sets of Latency Measurements .....                                          | 10                 |
| Figure 4 – IPDV Calculation.....                                                       | 11                 |
| Figure 5 – PDV Calculation.....                                                        | 11                 |
| Figure 6 – Example time series of Latency Measurement.....                             | 15                 |
| Figure 7 - PDF-CDF relationship .....                                                  | 15                 |
| Figure 8 – Conversion from Time Series to PDF to CDF to CCDF to Logarithmic-CCDF ..... | 16                 |
| Figure 9 –Time series latency data of Marked vs unmarked traffic .....                 | 16                 |
| Figure 10 – Probability Distribution Function (PDF) .....                              | 16                 |
| Figure 11 – Cumulative Distribution Function (CDF) .....                               | 17                 |
| Figure 12 – Complementary Cumulative Distribution Function (CCDF).....                 | 17                 |
| Figure 13 – CCDF on a Logarithmic scale.....                                           | 17                 |
| Figure 14 – Active Measurements .....                                                  | 18                 |
| Figure 15 – Using the TCP handshake to measure latency .....                           | 19                 |
| Figure 16 – TWAMP reference Model .....                                                | 21                 |
| Figure 17 – TWAMP Light reference Model.....                                           | 21                 |
| Figure 18 – STAMP reference Model .....                                                | 22                 |

## List of Tables

| <b>Title</b>                                       | <b>Page Number</b> |
|----------------------------------------------------|--------------------|
| Table 1 – Jitter definitions in the Industry ..... | 12                 |
| Table 2 – Understanding Latency Percentiles .....  | 14                 |

# 1. Introduction

Low latency is gaining importance in the internet experience. Low Latency is being approached as an end to end solution by operators. This includes Wi-Fi links in the home, DOCSIS links in the access network and core network segments. Providing lower end to end latency is a top priority for operators in the coming years. Measuring the latency in the network then becomes a vital requirement.

Operators (and 3rd party speed-test websites) have metrics on latency which they have reported and discussed with the community. Yet there is confusion surrounding the latency numbers and the ability to compare them between networks. The language and meaning of latency metrics (latency vs jitter, one-way vs round-trip, average vs 99<sup>th</sup> percentile), the latency measurement methods, what is being measured and when (peak vs off-peak periods), are varied. This paper provides clarity around these topics and discusses latency measurement architectures as well as best in class measurement tools to streamline latency measurement for the cable industry.

Operators want the ability to measure the difference in latency that is actually being delivered, before and after they deploy a new technology in their network, like DOCSIS 3.1 AQM, Low Latency DOCSIS, Low Latency WiFi etc. The latency portion of measurement reports (e.g. FCC's Measuring Broadband America initiative) are not optimal, and without a consistent measurement approach to latency, this could become a customer perception problem for the internet service providers. For new technologies that differentiate traffic, there are also questions around how latency for unmarked traffic vs marked traffic can be measured and reported. Operators will be asked to help troubleshoot latency issues and it will be important for them to identify latency within their networks vs. outside of their networks. This paper discusses the latency measurement frameworks which an MSO can integrate into their network deployment.

## 1.1. Quality of Experience

Latency is the time that it takes for a packet to make it across the network from a sender to a receiver and for the response to come back. Network latency is commonly measured as round-trip-time and is sometimes referred to as 'ping time'. As applications turn ever more interactive, network latency plays an increasingly important role for their performance. Applications that are real-time perform the best when latency is low, and adding more bandwidth without addressing latency doesn't improve the user experience. Packet forwarding latency can have a large impact on the user experience for a variety of network applications. The applications most commonly considered as latency-sensitive are real-time interactive applications such as voice over Internet protocol (VoIP), video conferencing such as Zoom, and networked online gaming. However, other applications are sensitive as well; for example, web browsing is surprisingly sensitive to latencies on the order of hundreds of milliseconds.

Test results in [ITU-T G.114], show that highly interactive tasks (e.g. speech, video conferencing and interactive data applications) can be affected by delays beyond 100 ms and users report significantly reduced mean opinion scores (MOS) when the voice delays are beyond the 150ms mark. The current [ITU-T G.114] recommends a maximum of a 150 ms one-way latency for VOIP applications.

Online games have some models [QoE and Latency] that indicate the impact that network parameters have on user experience. Some data exists to indicate that end-to-end round-trip latency should be kept below 25 ms or 50 ms in order to provide a good user experience, depending on the type of game (first person shooters, massively multiplayer online games, e-sports, etc.). When the operational response delay is less than 50 ms, the MOS scores tend to be high; when the operation response delay is around 100 ms,

the MOS decreases but is acceptable for some kinds of games, and when the operation response delay is beyond 200 ms, the interaction quality for the gamer is very poor.

If we assume that gaming servers centrally located in North America are serving gamers all over the continent, the round-trip time (RTT) on the fiber backhaul links for gamers in the west coast will be around 40 ms (assuming 4000 fiber kilometers between say San Francisco and Chicago, and speed of light in fiber as 0.67x speed of light in vacuum). These RTTs will be even higher for gamers across different continents, if they don't have separate gaming servers. So, for the games which require very low latency and latency variance, the 25 ms - 50 ms end-to-end target implies that the access network latencies need to be consistently in the order of 5 ms – 10 ms target to meet the requirements for online games.

Web browsing performance is traditionally tracked using page load time. Web content can be sourced from different servers and web browsers typically fetch resources from each server by opening up multiple TCP connections to the server. As there are multiple handshakes/interactions in each of the underlying protocols (DNS, TCP, TLS, HTTP) and all of those handshakes are impacted by the RTT, higher RTTs increase the page load time. See the paper [Belshe M] “More Bandwidth Doesn't Matter (much)” for experiments on how RTT affects page load time.

## **1.2. Latency in the Internet**

There are a few main contributors to the latency of a packet as it traverses the network. The switching/forwarding delay, propagation delay, serialization/encoding delay are some of the factors which affect packets as they go across various network devices and links, from the source to the destination. Queuing delay is usually the biggest contributor to latency, and is mainly caused by the current TCP protocol and its variants. This delay is encountered at the bottleneck links like the home Wi-Fi network or the access network. The majority of TCP implementations use loss-based congestion control, where TCP ramps up the number of bytes “in-flight” (i.e. its congestion window) until it sees packet loss, cuts its congestion window in half, and then starts ramping back up again until it sees the next packet loss. (When the buffers in the device transmit queues are full, a new arriving packet has to be discarded). This way TCP automatically adjusts its transmission rate to fully utilize the available capacity of the bottleneck link.

The result of this congestion window ramp-up and cut-in half mechanism is a saw-tooth behavior for the buffer going between partially full and totally full. In every home there are multiple users and applications that will use the same connection to connect to the internet. Applications other than TCP will suffer as the packets from those applications will arrive to nearly full buffer that may take tens or hundreds of milliseconds to drain. This can make web browsing perform poorly, and make VoIP, video chat, or online games unusable when other TCP based applications (e.g. streaming video) are in use.

## **1.3. Common techniques to reduce latency**

Setting the buffer sizes appropriately in each of the network devices is a first step to reduce the latency in the network. Active queue management (AQM) is the next step in mitigating queueing delay, where the basic idea is to detect the increasing queue created by TCP and, then drop a packet which will let TCP know to back off on its sending rate, much ahead of the time it takes to drop a packet when the buffer is completely full. There are variety of algorithms, such as random early detection (RED), Proportional Integral Controller Enhanced (PIE) etc., which an AQM system can implement.

The next stage in the evolution of latency reducing solutions is the dual-queue approach where the concept is to separate the traffic for queue-building applications from those applications/traffic flows which are non-queue building. See the paper [Greg W, SCTE 2019] Low Latency DOCSIS Overview and Performance Characteristics, for detail on these types of traffic flows and the dual queue approaches. Low

Latency DOCSIS and L4S technologies tackle the queueing delay by allowing non-queue-building applications to avoid waiting behind the full buffers caused by the current TCP or its variants.

## 2. View of latency measurements

Internet latency is crucial in providing reliable and efficient broadband services to customers who are connecting to servers across the country and the globe. The trend of real time gaming and other real time applications only accelerates the importance of accurately understanding the latency characteristics of the network. This bubbles up the task of latency measurement towards the top of an operator's priority list. Being able to accurately diagnose latency issues seen by residential or business customers is becoming more important. In order to support server selection in distributed /virtual computing environments measuring accurate latency becomes extremely important. Knowing the latency characteristics well allows an operator to make better decisions on which latency reducing technologies to deploy and where.

Accurately measuring network latency, however, is not an easy task due to lack of testing end points, lack of clock synchronization when needed, the sheer volume of collected data points, and aggregating and analyzing the data meaningfully. In addition, the time that latency is measured affects measurement results significantly due to network dynamics, volatile traffic conditions, and network failures.

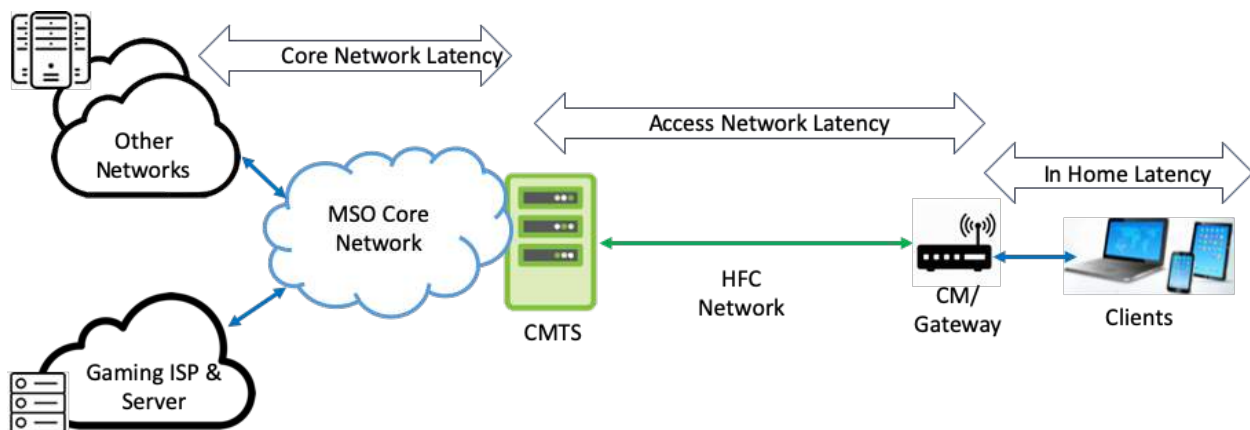


Figure 1 – MSO view of Latency Measurements

### 2.1. MSO Goals for Latency Measurement

Operators want to leverage existing available tools and standardized architectures to quickly set up a measurement infrastructure. Some of the common operator use cases and considerations are as follows.

- Identify Latency in 3rd party networks vs. MSO core network vs. Home network. In the Core network, there is a need to develop processes to identify routing issues, especially in the path to the egress point in the network which connects to a specific application server. For the Access & Home networks, it is extremely useful for an operator to be able to delineate latency from within the customer home (e.g. due to Wi-Fi) vs the access network latency vs the aggregation/core network.
- Operation Diagnostics Support: Operators would like to develop diagnostic tools, so that they can give meaningful information to their operations team. The use of latency measurements in NOC and field tester tools for live problem diagnosis is common at IP and Ethernet layers

- Operators would need to measure a variety of access and core architectures (R-PHY, FMA, Integrated) and need the measurement methods work across these range of deployments
- Network Architecture Analysis: The loss and delay performance metrics impact the scalability of the network and also on its behavior under load. For network architects, understanding both end-to-end network latency plus the contribution of the various links and nodes (network devices) that the network is comprised of is very useful.
- Understanding how to optimize the network deployments: e.g. with Distributed CCAP architectures an operator has to decide on a particular architecture, or where to place the physical or virtual components and decide on the location of certain functionality (e.g. MAC scheduler).
- There are many benchmarking purposes the latency measurement data can be used for e.g. different equipment (switches, routers) introduce different degrees of delay when processing packets. When moving from physical network elements to virtualized network elements, an operator needs to be able to quantify the latency difference.
- Lab latency measurements can compare the impact of introducing a new network element or configuration (e.g. a new technology like Low Latency DOCSIS) and verify the end user experience prior to deployment.
- Optimizing Network configuration: Appropriate latency measurement techniques can help diagnose intermittent issues (e.g. buffer overflows) and help fix them.
- Now with a goal of identifying per hop latency, operators need to identify the appropriate locations for the measurement end-points: end-device, gateway/CM, CMTS, router, interconnection point, etc.
- Any measurement architecture needs to support frictionless deployment of latency measurement infrastructure. This is dependent on how the specific measurement infrastructure is implemented and deployed (e.g. is it using hardware probes vs virtual probes). Scalability of the measurement platform across an entire operator becomes an important consideration.

## **2.2. Current National Latency Reports**

Broadband infrastructure is gaining the attention of various national communications regulators, as countries focus on enabling their people with high speed internet connectivity. As a part of this many of these regulators measure the broadband deployments and report on various metrics such as houses covered, speed tiers available etc. and also conduct network measurements on actual upload and download speeds. Latency measurements are now also becoming an integral part of these reports.

### **2.2.1. Measuring Broadband America**

In the United states, the Federal Communication Commission (FCC) runs the Measuring Broadband America (MBA) program. The MBA program is a nationwide study of consumer broadband performance and it collects network performance data from a representative sample of customers from each of the fixed Internet Service Providers (ISPs). See the paper [MBA FCC] ‘Ninth MBA Fixed broadband report’, for the latest speed and latency data reported. The MBA tests conducted are automated, direct measurements of the customers service during a single month and is done in collaboration with the measurement company SamKnows. Each volunteer customer connects a ‘Whitebox’ client device to their home network which performs the tests after finding the nearest test servers.

The MBA program measures latency by measuring the average round-trip time from the consumer’s home to the two closest measurement servers, one server chosen from each of two “pools” of servers. The report shows the median latency for each participating ISP and includes aggregated information for each ISP and type of access network. It reports the measured latencies for various DSL, cable and fiber based ISPs on an individual basis as well as aggregated. The MBA program has a limited number of test server

locations in each pool. Only six cities host test servers in both pools (an additional four cities host a server in only one pool). This means that client devices that are geographically distant from these six cities will report latency numbers that are more likely to be correlated to geography than to network capability. Difference in geographical distance to the server and also the distance in the of the number of hops internal and external to the ISP network, can make a difference in the number of network links the test packets have to travel across and ultimately the latency measured.

The MBA program latency and packet loss tests measure the round-trip times for approximately 2,000 packets per hour sent at randomly distributed intervals. Per the [MBA FCC] report, the latency and packet loss test records the number of packets sent each hour, the average round trip time and the total number of packets lost (a packet is considered lost if the packet's round-trip latency exceeds 3 seconds). The test computes the summarized minimum, maximum, standard deviation and mean from the lowest 99 percent of results. MBA determines the mean value over all the measurements for each individual's Whitebox and then computes a median value from the set of mean values for all the Whiteboxes.

### ***2.2.2. Measuring Broadband Canada Project***

The Canadian Radio-television and Telecommunications Commission (CRTC) has commissioned a study of the performance of broadband services sold to Canadian consumers. This project measures broadband Internet performance, including actual connection speeds, in Canadian homes. The CRTC collaborated with a number of Canadian Internet service providers (ISPs) and SamKnows, and produced a Measuring Broadband Canada Report, June 2020. See the paper at [MBC CRTC]. The report describes that, unlike in the US MBA program, the latency data was focused on Whiteboxes located within a 150km radius of the test server locations in order to minimize the effect of distance on measurements. See the paper [MBC CRTC] to understand the details on the average latency during peak hours for different Canadian service providers and access networks (Cable, DSL, Fiber) . Like the MBA report, the MBC report [MBC CRTC] also measures packet loss and average webpage loading times from a selection of websites.

### ***2.2.3. EU Broadband Report***

The European Commission has a vision around broadband connectivity and takes policy actions to turn Europe into a 'Gigabit Society' by 2025. In support of that the European Commission has commissioned a study to obtain reliable and accurate statistics of broadband performance across the different EU Member States and other countries.

### ***2.2.4. Speedtest (Ookla)***

Speedtest(Ookla) today publishes [SpeedTest] Market Reports as a guide to the state of fixed broadband and mobile networks around the world. Each report includes mainly speed (downstream and upstream) data and insights about country trends. Speed test data is based on the results of millions of tests run by Speedtest users. An individual user initiated Speedtest uses 'ping' to report the latency to the nearest Speedtest server. Speedtest is very relevant in the latency measurement landscape as that is how the majority of consumers understand what their service speed are and what latencies their connection achieves. Of course, consumers also tend to run Speedtests when they see an issue with the service or when they upgrade or get a new service, so this may also not be a representative sample across the consumers.




### 3. Latency Metrics

Each operator needs to track different metrics or KPIs when it comes to network latency. The network latency metrics important to operations teams will be different than what metrics are important to product or regulatory teams. Metrics can also be dependent on where the network is in the product life cycle. There are a variety of latency metrics to choose from and this section describes how to look at and understand latency.

As a packet travels across a network, the packet experiences different types of delays at intermediate hosts, routers, and network links. A host or a router needs time (processing/ switching/ forwarding delay) to process an incoming packet to determine its next hop. The packet also often waits in the transmit queue behind other packets (queuing delay). Transmission delay (serialization/encoding delay) is the time for a node to move out all the bits of the packet onto the link. Finally, it takes time for the packet to propagate over the link from one node to another. End-to-end latency is the sum of such delays at every step of the way.

#### 3.1. One-way Latency (or Packet Delay)

One-way latency is the total time it takes for a packet of data to travel from the sender to the receiver, across one or multiple hops. The one-way delay will be dependent on congestion of the network at the time the packet was sent. It will also depend on the topology of the network and the distance and routing decisions between the two end points. Measuring one-way latency also implies that the sender and receiver have synchronized clocks, which sometimes is a challenge to set up and maintain when the end points are across multiple network domains.



**Figure 2 – One way Latency vs. Round trip Latency**

#### 3.2. Round Trip Latency


Round trip time (RTT) or round-trip latency, is the time taken for a packet of information to travel from the sender to the receiver and back again. RTT is the total time it takes for a packet of data to travel from the sender to the receiver, across one or multiple hops, plus the total length of time it takes for receiver to send a packet back to the sender, through one or multiple hops.

The Round-trip latency is more often quoted, as it can be measured from a single point. It requires a process running on the other end to mirror the packet back. The RTT can vary if the return path is different from the forward path. The most common example for round-trip measurements is the ICMP Echo Request/Reply, used by the ping tool.

### 3.3. Singleton Measurements vs Sets of Measurements

A singleton measurement test can send one packet and calculate the one way or round-trip latency of that packet. That is not the most interesting as that is just one sample on the network which is carrying millions of packets. Latency varies with different factors such as the location of the two measurement end-points, and with time (due to changes in route selection or due to congestion). So, most latency measurement tests use multiple packets in a test for one way or round-trip measurements. This gives an operator a sample distribution of latency measurements and paints a better picture of the latency behavior. A test would measure the latency of each of the test packets, and then an operator could understand what the behavior of the network latency is across that set of packets. Having more data samples allows the operator to observe the variations and better understand the network latency in a way that correlates to what the customer will perceive and experience.

Operators typically run each of these tests multiple times a day to get a feel for the network latency variation over time. These sets of measurements could be performed over time for one user, these could be tests done across multiple users or it could be both: tests done over time and for multiple users.



**Figure 3 – Sets of Latency Measurements**

### 3.4. Jitter or Delay Variation

The term ‘jitter’ is a commonly used term to refer to variation in the latency of arriving packets over time. Though prevalent in the networking parlance, the term is considered deprecated by technical bodies like the IETF. The IETF [IETF RFC 5481] now uses the term "delay variation" for metrics that quantify a path's ability to transfer packets with consistent delay. Note that jitter can also be used to convey undesired variation in signals in contexts beyond IP packet transfer. (e.g. frequency or phase variations in electronic circuits in reference to a clock, or sampling jitter in analog-to-digital conversion of signals etc.)

The term jitter can be defined within a specific context in order to provide a meaningful metric for a specific application, or it is sometimes defined simply in a manner that is convenient to calculate. Most real-time voice and video applications employ a (de-jitter) buffer to smooth out delay variation encountered on the path. Many of the commonly used jitter definitions are aimed at helping designers of such systems choose the size of the de-jitter buffer.

[IETF RFC 5481] notes that various standards for delay variation have allowed flexibility to formulate the metric and so the specific formulations of delay variation must be well understood. All definitions of delay variation are derived from the one way or round-trip delay metrics. The networking industry has predominantly implemented two specific formulations of delay variation: Inter-Packet Delay Variation and Packet Delay Variation.

### 3.4.1. Inter-Packet Delay Variation, IPDV

A latency test or application will send a sequence of packets to measure one way or round-trip latencies. Inter packet delay variation (IPDV) is derived from such a sequence of latency measurements. It is simply the difference in latency of each packet as compared to the previous packet.




Figure 4 – IPDV Calculation

### 3.4.2. Packet delay variation, PDV

Packet Delay Variation, PDV, is also derived from a sequence of latency measurements where a single reference latency is chosen from the stream based on specific criteria. The most common criterion for the reference is the packet with the minimum delay in the sample. Other references such as average latency can be chosen as well. PDV is simply the difference in latency of each packet as compared to the one reference packet. In [ITU-T Y.1540] the ITU also chooses this definition of packet delay variation.




Figure 5 – PDV Calculation

### 3.4.3. Jitter Metrics in Use in industry

The formulations described in the previous sections result in a per-packet metric, which (across a set of packets) can then be summarized using descriptive statistics (e.g. mean, median, standard deviation, median absolute deviation, P99, P99.9, etc.) in order to come up with a summary of the delay variation across the set of packets.

There are different ways in which jitter definitions are used in different applications in the industry. [SamKnows], [Haste], [Excentis ByteBlower], [M-lab NDT] and [Network Next] use statistics derived from the PDV definition, while [WTFast], [3rdEchelon], [IETF RFC 3550] and [PingPlotter Pro] use statistics derived from the IPDV definition. The below table describes the way definition each of these entities use and how they aggregate it. As can be seen, there are significant differences in the meaning of the term from one entity to another.

**Table 1 – Jitter definitions in the Industry**

| Entity                                                 | Parameter                          | Definition                                                          |
|--------------------------------------------------------|------------------------------------|---------------------------------------------------------------------|
| <b>PDV based</b>                                       |                                    |                                                                     |
| SamKnows<br>(Network performance measurement platform) | Jitter                             | P99 PDV<br>(PDV referenced against min latency)                     |
| Excentis<br>(Byteblower traffic generator)             | Jitter                             | Standard deviation of PDV                                           |
| Haste<br>(Optimized routing for game traffic)          | Jitter                             | Standard deviation of PDV                                           |
| Network Next<br>(Optimized routing for online games)   | Jitter                             | jitter = 3* RMS(PDV)<br>(PDV referenced against min latency)        |
| MLab NDT<br>(Network test)                             | Jitter (round trip time variation) | max(PDV)<br>(PDV referenced against min latency)                    |
| <b>IPDV based</b>                                      |                                    |                                                                     |
| RTP protocol<br>(RFC3550)                              | Interarrival jitter                | Exponentially-weighted moving average of the absolute value of IPDV |
| PingPlotter Pro<br>(Ping statistics tool)              | Jitter                             | Average of the absolute value of IPDV                               |
| 3rdEchelon<br>(Internet Services company)              | Jitter                             | Average of the absolute value of IPDV                               |
| WTFast<br>(Gaming VPN solutions)                       | Jitter                             | Average of the absolute value of IPDV                               |

### 3.5. Descriptive statistics

Once we have a set of measurements (each of which is an individual latency measurement), a network operator wants to easily aggregate and make sense of those sets of measurements across the whole network and over time. The question is how best an operator can analyze the data to guarantee that the latency meets service requirements.

#### 3.5.1. Basic statistics

Many operators start with basic statistics like mean, median or min-max. Each of these numbers have their place, but for large populations of data they often hide the actual network behavior. Mean and median tend to hide outliers, especially the high latency events which may happen only during specific times. In contrast, the maximum is overly conservative and is easily distorted by a single outlier event.

- **Average:** The arithmetic mean or average, is the simply the sum of the set of the latency measurements divided by the number of measurements. The set of results of each experiment or an observational study can yield its own average number. For latency measurements, though the average maybe a starting point, it hides a lot of the variation in latency. Some of the much higher excursions are diluted by the mean, and thus averages hide high latency events which would ultimately impact the customer experience. Outliers also skew averages, so the average doesn't represent typical behavior either.
- **Min/ Max:** The maximum and minimum of a set of measurements are the largest and smallest value in the set of measurements. These are useful to understand the limits of the network. In the context of latency measurements one can separate lost packets as a separate measure, instead of considering it as infinite latency.

- **Standard Deviation:** The Standard Deviation is a measure of how spread out the latency measurement numbers are. The standard deviation is calculated as the square root of the variance (average of the squared differences from the Mean, for each sample). This gives a measure of the amount of variation or dispersion of a set of latency values. A low standard deviation indicates that the values tend to be close to the mean of the set, while a high standard deviation indicates that the values are spread out over a wider range.

### **3.5.2. Percentile Numbers**

Many descriptive statistics like mean, standard deviation, and skew are most meaningful when the underlying data follows a roughly normal (Gaussian) distribution. In contrast, even simple latency distributions are often heavily skewed with a set of values around a certain range, and with many fluctuations and outliers. As a result, these traditional statistics offer very little value in capturing or describing latency, but percentiles can generally be much more effective.

Percentiles allow a better understanding of the latency distributions than averages. A percentile is a value below which are a certain percentage of observations. Percentiles show the point at which a certain percentage of observed values occur. For example, the 95<sup>th</sup> percentile is the value which is just greater than 95% of the observed values, i.e. 95 percent of packets got a lower latency than the P95 value. For example, to obtain the 99<sup>th</sup> percentile of a collection of latency measurements from a network, an operator can sort them and discard the highest 1% of values. The largest remaining value is the 99<sup>th</sup> percentile. This value is the largest latency that will be seen for 99% of the measurements. An operator can choose a measure like the 90<sup>th</sup>, 95<sup>th</sup>, 99.9<sup>th</sup> (or even more nines) percentiles, which are typically denoted as P90, P95, P99 etc.

Network latencies between machines can be low when the network path is idle, but when there is significant network activity packets can take anywhere from a few milliseconds to hundreds of milliseconds, or even seconds. Since many network segments (particularly broadband links) are idle, or nearly so, for a significant portion of the day, the median latency and the minimum latency are often pretty close to one another. Long tail latencies occur when the higher percentiles begin to have values that are many times greater than the median. In a long tail latency distribution, the 99<sup>th</sup> percentile can be fifty times greater than the median value, much beyond normal distributions.

Percentiles are often used to find outliers. When a range of percentiles are computed one can estimate the data distribution more accurately. Another use for latency percentiles is to implement a threshold beyond which issues are flagged to the operator. An operator could also track a combination of a few different percentiles, such as P50, P75, P95, P99 and flag issues when any of them change significantly with respect to previous measurements or thresholds.

Now the question is which latency response time metric is more representative of the user experience. Is it the 95<sup>th</sup> percentile or the 99.9<sup>th</sup> percentile? The below table tries to show how to think about the impact to an application like gaming. Gaming traffic flows are typically 60 packets per second at a rate of 100kbps-200kbps in the upstream direction and 60 packets per second at a rate of 500kbps-1Mbps on the downstream. Gaming clients or servers expect packets to arrive at that consistent rate of 60 times per second and any packets which arrive with a much higher latency cannot be used and are essentially thrown away. As an example, 99% of the gaming packets have a latency of 40ms or less, while 1% of packets are delayed for anywhere from 100ms to 500ms. For or a real-time game this 1% ‘latency event’ happens (on average) once every 1.6 seconds and such network behavior is unwelcome in gaming environment and may be a showstopper in other applications. Based on this view, perhaps the P99.9 value would be a good starting point to represent user experience for online gaming.

**Table 2 – Understanding Latency Percentiles**

| Notation | Percentile Latency                           | Meaning                                      | Implication                                      | Impact for a gaming application  |
|----------|----------------------------------------------|----------------------------------------------|--------------------------------------------------|----------------------------------|
| P50      | 50 <sup>th</sup> percentile - median latency | 50% of packets got this latency or better    | 50 of 100 of packets got worse than this latency | Every other packet!              |
| P90      | 90 <sup>th</sup> percentile                  | 90% of packets got this latency or better    | 10 of 100 packets got worse than this latency    | 6 packets a second               |
| P95      | 95 <sup>th</sup> percentile                  | 95% of packets got this latency or better    | 5 out of 100 packets got worse than this latency | 3 packets a second               |
| P99      | 99 <sup>th</sup> percentile                  | 99% of packets got this latency or better    | 1 of 100 packets got worse than this latency     | 1 packet every 1.6 seconds       |
| P99.9    | 99.9 <sup>th</sup> percentile                | 99.9% of packets got this latency or better  | 1 of 1000 packets got worse than this latency    | 1 packet every 16.6 seconds      |
| P99.99   | 99.99 <sup>th</sup> percentile               | 99.99% of packets got this latency or better | 1 of 10,000 packets got worse than this latency  | 1 packet every 2 mins 46 seconds |

### 3.6. Histograms

A histogram is a graphical method for displaying the shape of a distribution and is particularly useful when there are a large number of observations. To construct a histogram, the range of values observed in the measurement is divided into a series of intervals or bins. The measurements are then classified into bins counting how many values fall into each interval. The bins are usually specified as consecutive, non-overlapping intervals of a variable. The bins/intervals are contiguous and are often of equal size.


The goal is to collect enough data points for good latency characterization. This means an operator needs to collect data to obtain acceptable precision for different percentile levels. A simple process in latency measurement is to record all the latency data over multiple sets of tests and then later analyze and sort the data into traditional histograms to get the required percentile data. Some alternatives to the traditional histograms with linear bins are logarithmic bins, or arbitrary bins. Linear bins in histograms require lots of storage to cover the range with good resolution, while logarithmic covers wide range of values but does not have good precision. Arbitrary bins work only when the operator already has a good feel for the interesting parts of the data range.

### 3.7. Visualization of Latency

Data visualization can reveal patterns and trends in the data, allows quick absorption of large amounts of data by network operators, and ultimately lets the operator understand the information and make decisions. This section describes some of the ways an operator can visualize latency data.

#### 3.7.1. Time series


A time series is a set of observations ( $x_t$ ) ordered in time, observed at a discrete set of (approximately) evenly spaced time intervals: at times  $t = 1, 2, \dots, N$ , where  $N$  is the length of the time series. The figure below, created using [PingPlotter Pro], shows a time series of ping data, once every second for 10 minutes. While the average ping time is ~12ms, one can quickly see that it is not the normal case and there are many latencies of 15-20ms and occasionally even up to 25 to 30ms.



**Figure 6 – Example time series of Latency Measurement**

### 3.7.2. Probability density function (PDF)

A Probability Density Function (PDF) is a statistical expression used in probability theory as a way of representing the range of possible values of a continuous random variable. For a continuous function, the probability density function (pdf) is the probability that the variable has the value  $x$ . The area under the curve represents the probability that variable will fall within an interval; and is expressed in terms of an integral between two points.  $\Pr[a \leq X \leq b] = \int_a^b f_X(x) dx$




**Figure 7 - PDF-CDF relationship**

### 3.7.3. Cumulative Distribution Function, CDF

The cumulative distribution function (CDF) of a random variable is another method to describe the distribution of random variables. A cumulative distribution function describes probability that a random variable takes on a value less than or equal to  $x$ . That is  $\Pr[X \leq x] = F_X(x)$

### 3.7.4. Complementary cumulative distribution function, CCDF


A complementary cumulative distribution function, answers the opposite question, i.e. how often is the random variable above a particular level  $x$ . To obtain the Cumulative Distribution Function (CDF), the integral of the PDF is computed. Then inverting the CDF results in the CCDF. (CCDF is the complement of the CDF or  $\text{CCDF} = 1 - \text{CDF}$ .) One can also plot the CCDF in a logarithmic scale so that the more interesting percentile values are easily discernible.



**Figure 8 – Conversion from Time Series to PDF to CDF to CCDF to Logarithmic-CCDF**


### 3.7.5. Example PDF/CDF/CCDF

Below is an example of some latency measurements performed in the lab, of round-trip times from a client to a server, which are separated by a Wi-Fi link and a DOCSIS link (CM and CMTS) with a pseudo Low Latency DOCSIS configuration. The time series figure below shows the latency measurements of unmarked traffic (in blue), while the latency of DSCP marked traffic is shown in orange. It also shows the various latency and jitter metrics of the unmarked traffic and how varied the numbers can be.



**Figure 9 –Time series latency data of Marked vs unmarked traffic**


The PDF figure below shows (using a histogram of 1ms bins) how different the two sets of latency measurements are, with the marked traffic flow(orange) having a lower and tighter latency numbers, while the unmarked traffic(blue) has latencies extending all the way from 100ms to 260ms.



**Figure 10 – Probability Distribution Function (PDF)**


The Cumulative Distribution function (CDF) figure below for the same data set, show the marked traffic flow(orange) having a lower P99 (~38ms), while the unmarked traffic has a higher P99 (~125 ms).






**Figure 11 – Cumulative Distribution Function (CDF)**

The CCDF figure below for the same data set, is essentially the same graph but inverted, with P99 readings closer to the bottom of the graph compared to the top.



**Figure 12 – Complementary Cumulative Distribution Function (CCDF)**

The logarithmic CCDF figure, is the same CCDF graph but drawn on a logarithmic scale for both axes. Here we can compare the P90, P99 or P99.9 and see the differences in the percentiles we are interested in clearly at this scale.



**Figure 13 – CCDF on a Logarithmic scale**


## 4. Latency Measurement architectures

### 4.1. Types of measurement

#### 4.1.1. Active measurements

Active measurements are conducted by generating traffic between two end points for the sole purpose of measuring the latency. For example, with ICMP ping a sender sends an ICMP packet(s) to the receiver, who replies back; the sender calculates the time between sending and receiving the packet(s). The measurement is considered to be active, as the reason the traffic is created and sent is to measure the latency between the end points.

Active monitoring involves injecting test traffic into the network, typically with the same forwarding criteria as the user traffic being monitored, and then measuring its performance. These tests can either be one-way (from site 'A' to site 'D' or round trip (from site 'A' to site 'D' and back to site 'A'), depending on what the operator wants to measure. Since the test traffic mimics the user traffic, active testing gives a packet by packet view of the end-to-end performance of a network with regards to such things as latency, delay variation, or packet loss.



**Figure 14 – Active Measurements**

Active testing can be performed between successively longer paths along the network route, for example, from site 'A' to site 'B' or site 'A' to site 'C'. With this the operator can segment the overall end-to-end path so that performance indicators can be derived on a per segment basis, giving visibility into where issues might be located. Active monitoring is the primary method for policing service level agreements, since it provides a real-time view of performance. Active monitoring requires two end points to be able to create test traffic and respond back to complete the measurement

#### 4.1.2. Passive measurements

Passive measurements are done simply by observing normal host-host interactions. Instead of measuring the latency of specially created test packets like in active measurements, passive measurements are based on the normal user packets that traverse the network. Passive measurements observe the traffic exchanged between two end-points and calculates the latency based on observed activity. For example, during normal interactions between host A and D, say during the initial handshake, a packet sent from A to D would be immediately responded by D as per the normal protocol interaction. If this transaction can be observed say at a location B, one can measure the time between sending the packet and receiving the response. Passive methods obtain similar measurements as an active measurement, without creating any new test traffic in the network, but are reliant on the presence of user traffic and can thus be skewed (for better or for worse) toward periods of time when more such traffic is present. Passive monitoring involves

capturing and analyzing live network traffic, or traffic statistics, at a specific point in the network, for example at the network interface to an application server, or at an aggregation router.

Passive monitoring does not require another host in the network to be involved in the process. Passive monitoring involves capturing some, or all, of the traffic flowing through a port for detailed, offline analysis of things like signaling protocols, application usage or top bandwidth consumers. Passive monitoring is suited for in depth traffic and protocol analysis, and can give visibility into the customers actual quality of experience.

#### 4.1.2.1. TCP Analysis

Analyzing the delay experienced by the TCP connection setup packets is an example of passive measurements. TCP uses a three-way handshake to establish a reliable connection. The TCP connection setup consists of a handshake with SYN, SYN+ACK, and ACK packets. The idea is to examine the data for outgoing connections, and compute the round-trip delay between the SYN & SYN+ACK packets as well as the SYN+ACK & ACK packet in the handshake. Since TCP connection endpoints normally respond immediately this is an easy way to compute the round-trip times.




Figure 15 – Using the TCP handshake to measure latency

## 4.2. Industry measurement architectures

This section describes some of the commonly used measurement architectures.

### 4.2.1. SamKnows Whitebox (dedicated test device solution)

SamKnows has developed a “Whitebox”, a dedicated device with a test suite, for measuring internet performance. These Whiteboxes are used by service providers, government regulators etc. and the tests can also be incorporated into network devices like modems or routers. The [SamKnows] test methodology includes many aspects of measuring consumer broadband performance: providing consumer volunteers with the Whiteboxes to run tests on consumer internet connections, the mechanism for collecting and aggregating the data, and finally the format for presenting the data.

The following describes the overall latency measurement methodology followed by SamKnows Whiteboxes. As described in [SamKnows] literature, upon start up, the Whitebox runs a brief latency measurement to all measurement servers hosted by an operator, or hosted by Samknows on their behalf. The server with the lowest round-trip latency is selected as the target for all subsequent measurements.

Below are some of the latency specific tests that the SamKnows Whitebox, or routers with the test functionality can run, as described in the [SamKnows] documentation.

- Latency and packet loss (UDP): This test measures RTT of small UDP packets between the Whitebox and a target test server. Each packet consists of an 8-byte sequence number and an 8-byte timestamp. The test operates continuously in the background and randomly distributes the sending of the packets over a fixed interval, typically 2000 samples per hour. It then records the number of packets sent, the average round trip time of these and the total number of packets lost. The test uses the 99<sup>th</sup> percentile when calculating the summarized minimum, maximum and average results on each Whitebox.
- Contiguous packet loss / disconnections (UDP): This test records instances when two or more consecutive packets are lost to the same test server. Alongside each event it records the timestamp, the number of packets lost and the duration of the event. By executing the test against multiple diverse servers, an operator can begin to observe server outages or and disconnections of the user's home connection.
- Latency, jitter and packet loss (Fixed rate UDP test): This test uses a fixed-rate stream of UDP traffic, a bi-directional 64kbps stream (representative of the G.711 voice codec), running between the client and test nodes. The standard configuration uses 500 packets upstream and 500 packets downstream. The server and client record the loss rate and the jitter observed. Jitter is calculated using the PDV approach described in [IETF RFC5481]. The 99<sup>th</sup> percentile is recorded and used in all calculations when deriving the PDV.
- Latency and packet loss (ICMP): This test measures the mean round trip time (RTT) of ICMP echo requests in microseconds from the Whitebox to a target test node.

#### **4.2.2. The M-Lab NDT (User initiated)**


M-Lab is a consortium of research, industry, and public-interest partners, and provides an ecosystem for the open, verifiable measurement of global network performance. All of the data collected by M-Lab's global measurement platform are made openly available, and all of the measurement tools hosted by M-Lab are open source.

M-Lab defines a test suite known as Network Diagnostic Tool (NDT), which is a single stream performance measurement of a connection's capacity for bulk transport (as defined in IETF's RFC 3148). NDT reports upload and download speeds, and latency metrics. The M-Lab NDT is run by users to test their internet connections. As described in [M-lab NDT], when the test is run, the client attempts to pick the nearest server from the geographically distributed network of servers provided by the M-Lab platform. The test suite uses a 10-second bulk transfer from the server to the client. The server is instrumented with the TCP kernel instrumentation and captures several variables of the TCP state machine every 5 ms of the test. NDT uses the TCP RTT samples as the latency data points and reports the difference between the minimum and maximum RTT observed during a test run.

#### **4.2.3. TWAMP**


Two-way measurements are common in IP networks, primarily because synchronization between local and remote clocks is unnecessary for round-trip delay, and measurement support at the remote end may be limited to a simple echo function. [IETF RFC 5357] specifies the Two-Way Active Measurement Protocol (TWAMP), which provides a common protocol for measuring two-way or round-trip measurement between network devices.

The [IETF RFC 5357] TWAMP defines a standard for measuring round-trip network performance between any two devices that support the TWAMP protocols. TWAMP consists of two inter-related protocols: TWAMP-Control and TWAMP-Test. The TWAMP-Control protocol is used to set up performance measurement sessions, i.e. to initiate, start, and stop test sessions. The TWAMP-Test protocol is used to send and receive performance-measurement probes i.e. exchange test packets between two TWAMP entities. The TWAMP measurement architecture is usually comprised of two hosts with specific roles, shown below. The first host (controller) consists of the control-client which sets up, starts, and stops TWAMP-Test sessions, and the session-sender which instantiates TWAMP-Test packets that are sent to the session-reflector. At the second host (responder) the session-reflector reflects the measurement packet upon receiving the TWAMP-Test packet. The responder can also have the TWAMP server that manages one or more TWAMP sessions.



**Figure 16 – TWAMP reference Model**

TWAMP Light is an alternative architecture which eliminates the need for the TWAMP-Control protocol and assumes that the Session-Reflector is configured and communicates its configuration with the Server through non-standard means. The Session-Reflector simply reflects the incoming packets back to the controller while copying the necessary information and generating sequence number and timestamp values. In TWAMP light, the roles of Control-Client, Server, and Session-Sender are implemented in one host (the controller), and the role of Session-Reflector is implemented in another host (the responder).




**Figure 17 – TWAMP Light reference Model**

TWAMP is more accurate than simple ping or traceroute measurements and is used by many operators in their transport, core and access networks. Several independent implementations of both TWAMP and TWAMP Light [IETF RFC5357] have been developed, deployed, and provide important operational performance measurements.

TWAMP is implemented in many of the core router products. TWAMP can provide accurate latency, jitter & packet drop KPIs, is supported by many probe vendors, and it can be integrated into network node equipment elements and CPE.

#### 4.2.4. Simple Two-Way Active Measurement Protocol

Simple Two-way Active Measurement Protocol (STAMP), is a newer IETF standard [IETF RFC 8762] which provides a simpler mechanism for active performance monitoring. It separates the control functions (vendor-specific configuration or orchestration) and test functions. STAMP also enables the measurement of both one-way and round-trip metrics (delay, delay variation, and packet loss)



**Figure 18 – STAMP reference Model**

### **4.3. Measurement considerations**

#### **4.3.1. Measurement under load vs quiet times**

Latency measurement tests need to ensure that testing is done over a variety of times to understand the variation between when the network is relatively lightly loaded and peak load time. This can also be used to measure self-congestion vs. network-congestion. For example, measuring latencies at peak time in the evening when most of the subscribers on the plant are online is likely to catch incidents when the network is congested due to high overall network load. Another way latency numbers can be affected is the load within a single user's home itself or even within the same client. If multiple devices in the home are using the network for various purposes like consuming video, voice calls, gaming, then a latency measurement will likely yield different numbers than running the test at quiet times. The path to various servers can change automatically to accommodate network/routing changes, so measured latency may vary over time and it may be appropriate to get a broad picture of latency including such situations.

#### **4.3.2. Window over which the measurement is done.**

Every latency measurement test can have a different purpose; one could be for a quick and immediate diagnostic purpose, while another could be to gather long term statistics. To diagnose issues in the network, an operator will need to consider the correct amount of time to run a test, how many latency samples will be collected in each run, and how often the test will be run. This could include sample rates of once per hour, once per minute, once per second, and as frequently as 50 times per second. The sampling rate and the number of measurements run will depend on the ultimate goal of the operator. If the goal is to reflect the worst gaming experience, then more measurements which mimic the game traffic flows will give us a better idea of the performance of the network.

To understand latency, one has to consider the entire distribution of latency measurements. While it is important for operators to look at latency numbers at the 99.9<sup>th</sup> percentile or higher, many monitoring systems stop at the 90<sup>th</sup> or 95<sup>th</sup> percentile. The reason is simply because it requires larger amounts of data to be collected, stored and analyzed. The data collected by most monitoring systems is usually summarized in small, five or ten second windows. Given we can't meaningfully average percentiles or derive five nines from a collection of small samples of percentiles, there is no way to confidently know what the 99.99<sup>th</sup> percentile for the minute or hour was. A related question is how many total samples are needed to get valid statistics. If an operator wants to measure the 99.9<sup>th</sup> percentile latency, then at least 1000 latency measurements are required, and a lot more (at least 2000-8000) would be needed to have an accurate statistical estimate.

### **4.3.3. Off-net and On-net testing**

Active measurement architectures (e.g. SamKnows) may use client devices which run bandwidth and latency tests to a specific measurement server. A majority of test servers used by SamKnows customers are off-net, i.e. hosted on the internet outside the operator network. Reporting results to target servers off an ISP's own network represents a 'real world' experience for end users. However, an ISP is not in control of the paths that get to the server and would also like to understand and debug issues within their own network. Hence many ISPs install test servers inside their network ("on-net") to allow them to segregate on-net and off-net performance.

With both on-net and off-net servers in use, operators can see the difference in performance internal to their network vs. external to it. The results can be used to troubleshoot peering links, routing issues, or simply rule out any capacity problems within the operator's own network. Consequently, any active measurement deployment should have a mix of on-net and off-net servers.

### **4.3.4. Latency Measurement Test definitions**

When designing latency measurement tests, an operator needs to define the test and the associated parameters such as the test traffic flow (i.e. the packet size and rate used), whether to test under load or without load, and the periodicity of the measurements.

Many IP network switches and routers need the full packet to be clocked into the device before it can be forwarded to the next networking device in the path to the end destination. This delay is referred to as a serialization delay and these delays are often tested using 64-byte packets. For example, a 64-byte packet will have serialization delays of 5.12  $\mu$ sec when clocked in using a 100 Mbps port. However, serialization delays are usually proportional to the size of the packet. If the size of the packet was 1280 bytes, the serialization delays would be twenty times bigger at 102.4 microseconds. Though this doesn't include the processing delays through a device (router, switch, CMTS, CM), it gives a sense of the interaction between packet size and link speed (interface bandwidth) that each node in the network could add as an absolute minimum

Small (say 64 byte) UDP packets sent every few seconds from a test node is a good place to start for RTT measurements. Latency Tests which mimic the gaming experience, (e.g. 150 Kbps upstream, 600 Kbps downstream, ~200-byte packets) would be a good data set to collect to understand the impact to gaming or other real-time audio-conferencing services. Latency tests with bigger size packets (1500 bytes) could also be used to expose any packet size limitations in the network.

When testing latency, it is also a good idea to understand the latency when the network is under load vs. when it is not. Latency testing with load is typically done by running both a downstream and upstream speed test or something equivalent at the same time as doing latency measurements. While the speed test is running, the latency under load test can send packets to a target server and measures the round-trip time and number of packets lost. The test packets should be sent equally spaced over the duration of the speed test.

### **4.3.5. Marked traffic vs Unmarked traffic.**

Differentiated services or DiffServ [IETF RFC 2474] specifies a simple mechanism for classifying and managing network traffic and providing quality of service (QoS) on modern IP networks. DiffServ can, for example, be used to provide low-latency to critical network traffic such as voice or streaming media while providing simple best-effort service to non-critical services such as web traffic or file transfers.

The six most significant bits of the DiffServ field (previously ‘type of service’ (TOS) field) in the IP header are called as the DSCP (differentiated services Code point) and the last two bits are the Explicit Congestion Notification (ECN) field. Routers at the edge of the network classify packets and mark them with their DSCP value in a Diffserv network. Other network devices in the core that support Diffserv use the DSCP value in the IP header to select a per hop behavior for the packet and provide the appropriate QoS treatment. Various applications and services (typically UDP based) can also mark the traffic they generate with specific DSCP values. For example, the popular video conferencing application Zoom uses a default DSCP marking values of 56 for audio, 40 for video, and 40 for signaling.

In the Low Latency DOCSIS technology, by default, the traffic within an Aggregate Service Flow is segmented into the two constituent Service Flows by a set of packet classifiers that examine the DSCP field and the ECN field. Specifically, packets with a Non-Queue Building DiffServ value, 0x2A, per a current [IETF NQB PHB] draft, will get mapped to the Low Latency Service Flow, and the rest of the traffic will get mapped to the Classic Service Flow.

In the context of Low Latency DOCSIS and other technologies which support dual queue mechanisms, the question is how can we modify latency measurement tests to also report metrics on unmarked traffic as well as marked traffic. One solution is to run any test twice, once as currently designed without any packet marking, and once with marked DSCP packets, and report results on both. As more games and other applications start marking their packets, public internet measurement reports will also have to start reporting latencies on both types of traffic.

## 5. Conclusion

Interactive applications like gaming or real-time communication, where real-time responsiveness is required, are more sensitive to latency than bandwidth. These applications really stand to benefit from technology that can deliver consistent low latency. Operators need to understand the latency characteristics of their network and be able to delineate the latencies seen in the customer home, the access network, and the MSO-core network. Using a common set of metrics to describe latency is the first step in understanding the state of the networks. Round trip times are relatively easy to collect compared to one-way latencies. Multiple sets of measurements paint a better picture of the latency characteristics than single measurement. Using averages to measure latencies can be misleading, so an operator can choose better performance indicators such as the 99<sup>th</sup> or 99.9<sup>th</sup> percentile to track and understand latency behavior over time. Latency is being measured by national entities, raising the importance of operators to have their own latency measurement infrastructures. Active measurement techniques give an operator good control over the testing and a better understanding of the network over various times and conditions.

## Abbreviations

|      |                                 |
|------|---------------------------------|
| bps  | bits per second                 |
| ms   | millisecond                     |
| RTT  | Round trip time                 |
| TCP  | Transmission Control Protocol   |
| UDP  | User Datagram Protocol          |
| IETF | Internet Engineering Task Force |
| RMS  | Root Mean Square                |



# Bibliography & References

- [ITU-T G.114] *One-way transmission time* <https://www.itu.int/rec/T-REC-G.114> , Recommendation G.114 (05/03) & Annex B ITU-T G.114 (05/2000)
- [QoE and Latency] Saldana J., Suznjevic M. (2015) QoE and Latency Issues in Networked Games. Handbook of Digital Games and Entertainment Technologies. [https://doi.org/10.1007/978-981-4560-52-8\\_23-1](https://doi.org/10.1007/978-981-4560-52-8_23-1)
- [BelsheM] “More Bandwidth Doesn’t Matter (much)”: Mike Belshe, Google <https://www.belshe.com/2010/05/24/more-bandwidth-doesnt-matter-much/>
- [Greg W, SCTE 2019] *Low Latency DOCSIS Overview And Performance Characteristics, SCTE 2019* , Greg White, Karthik Sundaresan, Bob Briscoe
- [MBA FCC] Ninth Measuring Broadband America Fixed Broadband Report <https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-broadband-america-program-fixed>
- [MBC CRTC] Measuring Broadband Canada Report <https://crtc.gc.ca/eng/publications/reports/rp200601/rp200601.htm>
- [EU Broadband] European Commission Broadband Connectivity <https://ec.europa.eu/digital-single-market/en/connectivity>
- [SpeedTest] Speed Test reports <https://www.speedtest.net/global-index/united-states#fixed>
- [ITU-T Y.1540] *Recommendation ITU-T Y.1540, 2019, Internet protocol data communication service –IP packet transfer and availability performance parameters*
- [IETF RFC 5481] *Packet Delay Variation Applicability Statement*
- [IETF RFC 3550] *RTP: A Transport Protocol for Real-Time Applications*
- [3rdEchelon] <http://www.3rdechelon.net/jittercalc.asp>
- [Haste] <https://haste.net/2017/08/23/what-is-jitter/>
- [Excentis ByteBlower] <https://www.excentis.com/products/byteblower>
- [Network Next] <https://www.networknext.com>
- [WTFast] <https://www.wtfast.com/en/>
- [M-lab NDT] <https://www.measurementlab.net/tests/ndt/>
- [SamKnows] <https://samknows.com> , <https://samknows.com/technology/tests/latency-loss-and-jitter#latency-jitter-and-packet-loss-udp>
- [PingPlotter Pro] <https://www.pingman.com> , <https://www.pingman.com/kb/article/what-is-jitter-57.html>
- [IETF RFC 5337] *A Two-Way Active Measurement Protocol (TWAMP)*
- [IETF NQB PHB] *A Non-Queue-Building Per-Hop Behavior (NQB PHB) for Differentiated Services, draft-ietf-tsvwg-nqb-01, G.White*

# **Smart Gateways**

## **Active A.I. in Subscriber Networks**

A Technical Paper prepared for SCTE•ISBE by

**Kyle Haefner**

Senior Security Engineer

CableLabs

858 Coal Creek Cir, Louisville, CO 80027

303-661-3803

k.haefner@cablelabs.com

# Table of Contents

| Title                                    | Page Number |
|------------------------------------------|-------------|
| 1. Introduction.....                     | 3           |
| 2. Background.....                       | 3           |
| 2.1. Supervised .....                    | 4           |
| 2.2. Unsupervised .....                  | 4           |
| 3. Methodology .....                     | 5           |
| 3.1. Data-Sets and Data Collection ..... | 5           |
| 3.1.1. Datasets .....                    | 5           |
| 3.1.2. Attack Data .....                 | 5           |
| 3.2. Data collection.....                | 6           |
| 3.3. Device Complexity.....              | 6           |
| 3.4. Results Methodology .....           | 8           |
| 4. Results .....                         | 9           |
| 4.1. Complexity Results.....             | 9           |
| 4.2. Behavior Results .....              | 10          |
| 4.3. Overall Results .....               | 12          |
| 5. Conclusion .....                      | 13          |
| Bibliography & References .....          | 14          |

## List of Figures

| Title                                               | Page Number                  |
|-----------------------------------------------------|------------------------------|
| Figure 1: Data Collection Architecture.....         | 6                            |
| Figure 2: Spectrum of complexity.....               | 7                            |
| Figure 3: Highest NSR Roku Express .....            | 9                            |
| Figure 4: Median NSR Android .....                  | 9                            |
| Figure 5: Lowest NSR Eufy Light .....               | 10                           |
| Figure 6: Home Device NSR Complexity.....           | 10                           |
| Figure 7: Device Behavior Boundaries .....          | 11                           |
| Figure 8: F1 Score vs NSR Complexity [Home].....    | 12                           |
| Figure 9: F1 Score vs NSR Complexity [Lab] .....    | 13                           |
| Figure 10: F1 Score vs NSR Complexity [UNSW] .....  | 13                           |
| Figure 11: F1 Score vs NSR Complexity [SCADA] ..... | Error! Bookmark not defined. |

## List of Tables

| Title                                           | Page Number |
|-------------------------------------------------|-------------|
| Table 1: Malware Data.....                      | 5           |
| Table 2: Features from network flows .....      | 6           |
| Table 3: Confusion Matrix for IoT Traffic ..... | 8           |
| Table 4: NSR and F1 Scores .....                | 12          |

# 1. Introduction

In the last several years progress toward securing Internet of Things (IoT) devices has been made on several fronts. There are now mature specifications for IoT devices that require with encryption, authentication and authorization for every device (1). Governments and industry have released baselines (2), (3), (4) that provide guidance on what should constitute a secure device. There is even recent legislation at the state level aimed at enforcing security in IoT (5).

None of this will guarantee secure devices. There will *always* be devices that are exposed, unpatched and vulnerable. Even companies and manufacturers that prioritize security will inevitably find themselves with vulnerabilities inherited in the supply chain from decades old code like Ripple20 (6). Combine this with malware like Mirai that is constantly being updated to take advantage of these newly discovered vulnerabilities (7) and it becomes clear that building strong security into individual devices is simply not enough. The question that now needs to be answered is, can secure systems be built from networks of potentially insecure devices?

The question posed above is not a mere hypothetical one. Today's subscriber networks consist of not just a heterogenous mix of devices, but also the implicit mix of vulnerabilities and attack surfaces inherent in today's complex home networks. To address this problem in a comprehensive and systematic way, intelligence must be added to the network so as to give the network the ability to know the devices running on it, learn how those devices behave and be capable of actively and surgically blocking traffic that is outside the bounds of what is deemed normal.

This research presents a method whereby a centralized router/gateway can learn a device's behavior on the network and based on that behavior, determine normal and abnormal behavior from that device. The model presented in this paper takes advantage of the predictability of an IoT device's network footprint by developing a formalized measurement of complexity for each device. Low complex and simple devices are more accurately modeled and thus can be more confidently managed autonomously by the network.

After describing the framework necessary to measure the complexity of network devices, this work then uses this complexity measure to inform and tune an anomaly detection algorithm to construct a behavioral model for each device. This tuned model represents the behavior footprint of each device learned from its network traffic and forms the basis for differentiating normal traffic from abnormal.

To demonstrate the efficacy of this model, this work analyzes boundary of each device's learned behavior against seven common types of malware traffic from infected IoT devices. Finally, to illustrate that the model can be effectively applied to a broad spectrum of devices, four different IoT datasets were analyzed: one residential dataset, two lab datasets, and a dataset based on commercial IoT devices. The results show that this model can be an effective way to actively block Distributed Denial of Service (DDoS) attacks and malware traffic especially on low complex devices.

## 2. Background

There has been a great deal of research in the past few years to grant the network with the capability to learn what devices are running on it, (8) and how to automatically block devices from contributing in DDoS attacks. The majority of recent academic research in IoT behavior and fingerprinting relies on various machine learning (ML) techniques. The ML techniques can be broadly categorized into two main groups, supervised and unsupervised.

## 2.1. Supervised

Supervised learning requires a large corpus of labeled data. The work that takes advantage of supervised learning typically tries to classify a device on the network based on previous traffic that has been labeled as that same device.

Loepz-Martin et al. (9) build a network traffic classifier (NTC) using a recurrent neural network (RNN) and apply it to labeled IoT traffic. The goal of this is to identify the types of traffic and services exhibited by an IoT device as a step toward identifying the device.

Miettinen et al. (10) have developed a method, called IoT Sentinel, that uses machine learning to designate a device type on the network, referred to by the authors as a device fingerprint. Using the random forest algorithm and 23 network features they were able to identify device types on the network based on the device's traffic. The 23 features are based on layer two, three and four of the OSI networking stack. Expecting that the body of the packet will be encrypted, all the features the authors employed are based on unencrypted parts of the traffic like IP headers information.

Bezawada et al. (8) build on the work done in Miettinen using a machine learning approach to broadly identify the device and place it in a predefined category, such as a light bulb. According to the authors, even devices from different manufacturers can be placed into general categories such as two separate light bulbs can be identified and placed into a lighting category and addressed by security policies based on this category.

Supervised methods are generally highly accurate but require large examples of labeled traffic to adequately learn. These methods generally cannot classify things that were not present in the learning dataset and are unable to classify new or unknown devices.

## 2.2. Unsupervised

Unsupervised learning does not need data with labels and instead tries to learn underlying patterns in the data itself. It has various advantages in the context of IoT security as there will often not be labeled data available and there will always be new devices for which there exists no labeled data.

AuDI (11) implemented an autonomous device-type identification that uses the periodicity of device communications resulting in abstract device categories that could be used to enforce access control policies. DioT (12) extends the AuDI classification model to create a federated approach by aggregating device anomaly detection profiles.

Ortiz et al. (13) set up a probabilistic framework to monitor device behavior using a Long-Term Short-Term Memory (LSTM) neural network, to learn from inherent sequencing of TCP flows to automatically learn features from device traffic with the intent of categorizing devices and distinguishing between IoT devices and Non-IoT devices. The authors are able to identify previously known devices after only 18 TCP-flow samples and categorize devices into two classes IoT and Non-IoT

### 3. Methodology

#### 3.1. Data-Sets and Data Collection

##### 3.1.1. Datasets

This work analyzed four different datasets of IoT traffic:

- **home**: sourced from a residence and represents normal usage traffic contains 37 days of traffic from over 25 devices that are a mix of typical IoT devices and more general devices such as laptops and smartphones,
- **lab**: traffic is from several devices in a lab at Colorado State University, it contains approximately 22 devices and several month's worth of traffic,
- **unsw**: IoT traffic from a lab at University of New South Wales (UNSW) contains nearly three weeks of traffic from 29 devices that are mixed IoT and general-purpose devices,
- **scada**: traffic from a Supervisory Control and Data Acquisition (SCADA) test network at Colorado State University's Methane Emissions Test and Evaluation Center. It has nearly 40 devices mostly made up of SCADA devices called Lab jacks.

##### 3.1.2. Attack Data


The malware data comes from Stratosphere Laboratory (14) and contains 20 captures where malware was executed on IoT devices, and 3 captures for benign IoT devices traffic. This dataset was first published in January 2020, with captures ranging from 2018 to 2019.

**Table 1: Malware Data**

| Attack Name             | Description                                                                               | Unique Flows |
|-------------------------|-------------------------------------------------------------------------------------------|--------------|
| C&C                     | This is traffic where a device is connecting to a remote command and control server.      | 30           |
| C&C Heartbeat           | This is traffic that is meant to monitor the status of an infected host.                  | 3            |
| C&C Torri               | This is command and control traffic specifically from the Torri botnet                    | 1            |
| C&C FileDownload        | This is traffic from an infected device downloading a file or malicious payload           | 7            |
| DDoS                    | This is traffic where a device is participating in a distributed denial of service attack | 1            |
| Part of Horizontal Scan | This is traffic where a device is scanning locally on the network                         | 49959        |
| Okiru                   | This is traffic specifically from the Okiru botnet.                                       | 99888        |

### 3.2. Data collection

For the home, and lab datasets data was collected in the form of IPFIX/Netflow flows from a centralized router to a flow collector as shown in Figure 1: Data Collection Architecture. Flows were saved in a SQL database. For the UNSW data set and the SCADA data set, the pcap files were analyzed by an open source tool called Joy (15). Joy transforms pcap files into json data which was then loaded into a SQL database.




**Figure 1: Data Collection Architecture**

**Table 2: Features from network flows**

| Feature                  | Description                                                   |
|--------------------------|---------------------------------------------------------------|
| IPv4 Source Address      | IP Address of the device                                      |
| IPv4 Destination Address | IP Address the device is connecting with on the network       |
| In Packets               | Number of packets received by the device                      |
| Out Packets              | Number of packets sent by device                              |
| L4 Source Port           | Source port on device                                         |
| L4 Destination Part      | Remote port that the device is connecting with on the network |
| Protocol                 | IANA protocol, eg UDP,TCP                                     |

### 3.3. Device Complexity

The complexity of a device on the network is based on its network traffic. Sensors that talk to a relatively few network endpoints will be measured as low complexity devices. This is compared devices such as laptops and smartphones that make highly varied requests to many network endpoints. These devices will be measured as high complexity devices. To illustrate this, Figure 2 shows where devices fall in a range from low complexity to high complexity.



**Figure 2: Spectrum of complexity**

### General Purpose Device

A device that is capable of running multiple user-space applications. Some examples are smart-phones, tablets, laptops, some streaming devices, and smart TVs. These devices will have higher network complexity.

### Single Purpose Device

A device that generally runs a single application. They often are capable of only one or two threads. These devices will have lower network complexity.

For a network flow there are several features that can be examined as shown in Table 2. This research focused on destination-based features. The complexity analysis is based only on destination IP address and destination port, though other features could be added.

To measure a device's complexity on the network this work uses a concept similar to one used in communications, called the Signal to Noise Ratio (SNR) which compares the level of desired signal to the background noise. Within the context of Internet traffic from a device, the signal is defined as traffic data points that can be clustered, and the noise is defined as the datapoints that cannot be clustered. Some of the devices measured had no noise components making the SNR ratio undefined. To account for this this research uses the reciprocal if the SNR, referred here as the noise to signal ratio (NSR) shown in Equation 1. In the case where there are zero points of noise the NSR=0.

This measure of complexity uses the DBSCAN (16) clustering algorithm to compute the number of clusters (signals) and the non-clusters (noise). This algorithm is good at finding areas of high density that are separated by areas of low density. The DBSCAN algorithm has several advantages in that it can find clusters of arbitrary shape and size, and can include clusters that are non-convex, unlike other clustering algorithms such as k-means.

The DBSCAN algorithm takes in two primary parameters, a distance parameter  $\epsilon$  and a number of points that are within that distance called, *min\_samples*, to form a cluster. To automatically find the distance parameter this work uses the device's *IP\_Spread* as shown in Equation 1: Distance Calculation. The *IP\_Spread* is the set of IP addresses where the first order octet is unique and represents the total number of unique network that the device connects to.

### Equation 1: Distance Calculation

$$\epsilon = 128 * IP\_Spread$$



128 is the midpoint of the address space of a class C network. Experimentally setting  $\text{min\_samples} = 10$  was a good starting value for the total number of neighborhood points necessary for a point to be calculated as a core point.

The number of clusters found by the DBSCAN algorithm and the number of non-clusters is used to calculate the NSR for the device using Equation 2.

#### Equation 2: Noise to Signal Ratio

$$NSR = \frac{n_{noise}}{n_{clusters}}$$

### 3.4. Results Methodology

To evaluate the results of the model a dataset consisting of traffic from several different malware types was used to see how well the model was able to predict normal traffic from abnormal attack traffic. Evaluating the model produces four metrics on how the model is performing. These four metrics are:

- true positives (TP): malware traffic correctly identified as malware traffic,
- true negatives (TN): normal traffic correctly identified as normal traffic,
- false positives (FP): normal traffic incorrectly identified as malware traffic,
- false negatives (FN): malware traffic incorrectly identified as normal traffic.

Table 3 shows the confusion matrix for this analysis.

**Table 3: Confusion Matrix for IoT Traffic**

|                         | Predicted: Normal Traffic | Predicted: Malware Traffic |
|-------------------------|---------------------------|----------------------------|
| Actual: Normal Traffic  | TN                        | FP                         |
| Actual: Malware Traffic | FN                        | TP                         |

From these four metrics there are two important factors that are often used, precision, the ratio of correct positive predictions to the total predicted positives and recall, the ratio of correct positive predictions to the total positive examples.

#### Equation 3: Precision

$$P = \frac{TP}{TP + FP}$$

#### Equation 4: Recall

$$R = \frac{TP}{TP + FN}$$

And finally, the balanced method for measuring the efficacy of a prediction model is called the F1 score. The F1 score is the harmonic mean of the precision and recall.

### Equation 5: F1 score

$$F_1 = 2 * \frac{P * R}{P + R}$$

In Section 4 the results are shown how each device model performs against the malware traffic using the F1 score. An F1 score of 1 is a perfect score, it means that the device model correctly identified all of the device traffic as normal and all of the malware traffic as abnormal.

## 4. Results

The results across all four of the datasets support the hypothesis that devices can be measured for complexity based on their network flows, and that this measurement can be used to categorize devices into two separate groups, single purpose devices and general-purpose devices. Further, the results support that measuring the complexity of a device is relevant to improve modeling of devices. For brevity only three devices from the Home dataset are shown, one of high complexity, one of median complexity and one of low complexity. Figure 8, Figure 9, Figure 10, and Figure 11 show the results of all devices for each data set analyzed and averaged against the seven malware traffic types.

### 4.1. Complexity Results

Figure 3 shows the Roku Express device in the home dataset with the largest NSR value of 52.25. Only four clusters were found with 209 points of noise. This is compared to the J Android device which falls in the middle of all devices with an NSR of 10.2 as calculated from 19 clusters and 194 points of noise shown in Figure 4. Finally, in Figure 5 we see the lowest complex device on the home network which is the Eufy light bulb. This device has only 2 clusters and zero points of noise leading to an NSR value of 0.




Figure 3: Highest NSR Roku Express






Figure 4: Median NSR Android



**Figure 5: Lowest NSR Eufy Light**

Figure 6 shows the NSR complexity of all the devices on the home network. A line is drawn at the average NSR value of 8.5. Almost all the devices that have an NSR above this value are general purpose devices and conversely almost all below 8.5 are single purpose devices. Streaming devices such as Roku and Smart TVs were consistently measured across the four datasets as higher complexity devices. This is likely because these devices are often Linux OS based devices capable of running several user space applications.



**Figure 6: Home Device NSR Complexity**

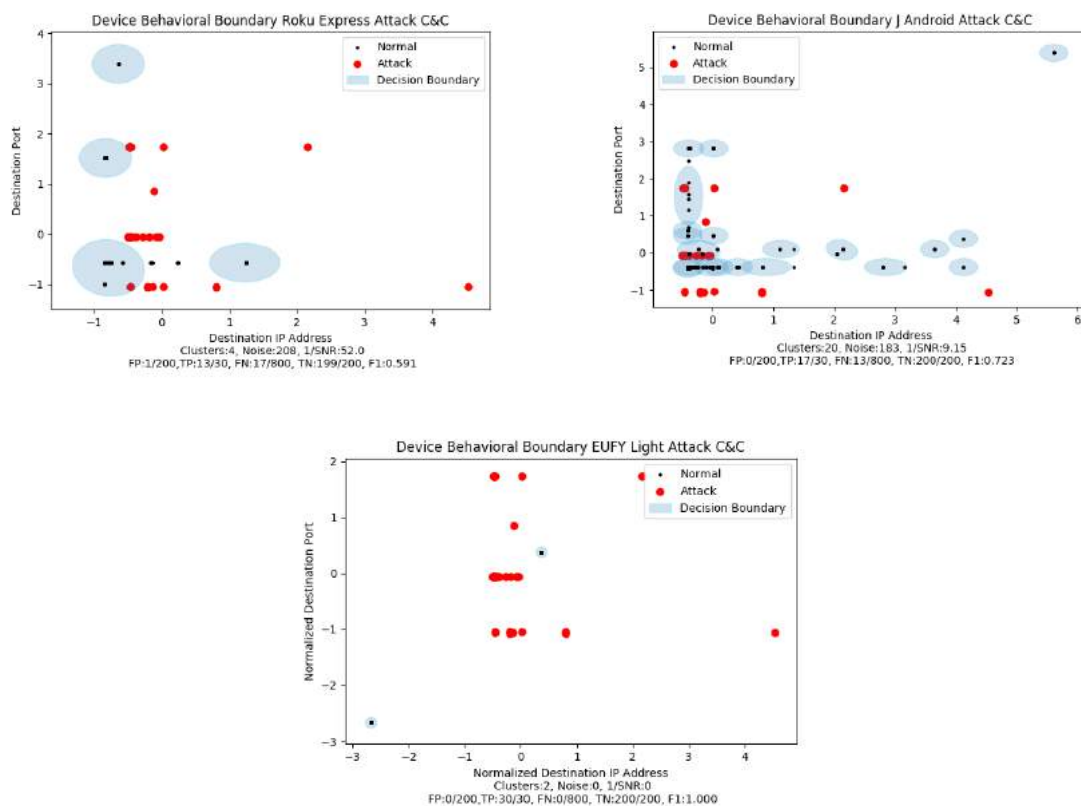
## 4.2. Behavior Results

Figure 7 shows the three devices, the Roku Express, J Android, and the Eufy light bulb. Each figure shows the normal traffic as black dots, the learned behavior boundary as blue ellipsoids generated from the Gaussians, and the attack traffic represented as red dots. In each figure the port and IP address have been normalized. Normalizing both improves classification and makes it possible to show the traffic and boundaries in non-logarithmic space.

The Roku Express has the highest NSR value of 52 in the home data set with only four clusters and 208 points of noise. This leads to a relatively poor model for the device with a 0.5% false positive rate against the C&C malware traffic. Visually, the figure shows large ellipses that fail to contain several of the normal traffic within their boundaries leading to 17 false negatives. Additionally, some of the ellipses overlap with the malware traffic leading to false positive identification. The number of true positive outliers detected is also poor with the model only detecting 43% of the malware traffic in the C&C traffic. This gives the model an F1 score of 0.59.

For the J. Chromebook device with an NSR value of 8.6 which is just slightly higher than the data set average of 8.548. This NSR value comes from a total of 15 clusters and 129 points of noise and leads the GMM to find a total of 15 ellipsoids. As compared to the model of the Roku, the J. Chromebook shows smaller margins on the ellipsoids that are centered around the normal traffic. The false positive rate of the J. Chromebook model is double that of the Roku model at 1%, however, there is less overlap in these ellipsoidal boundaries with the attack traffic than in the Roku model leading to a true positive accuracy of 100% for the malware traffic. This in turn leads to a much higher F1 score of 0.968 against the C&C malware.

Finally, in the figure for the Eufy light bulb there are only two ellipsoids generated by the model. The Eufy device is a light bulb and has the lowest measured NSR complexity in the home dataset with an NSR of 0 that is composed of just 2 clusters and 0 points of noise. This leads to a behavioral boundary tightly coupled around the normal traffic, with a 0.0% false positivity rate and a 100% accuracy in identifying the malware traffic. This leads to a perfect F1 score of 1.00.



**Figure 7: Device Behavior Boundaries**


### 4.3. Overall Results

Every device in each data set was modeled and evaluated against each of the seven malware traffic types. The average NSR complexity score, and the average F1 score for each of the four datasets is shown in Table 4. The table shows the inverse relationship between average NSR complexity and average F1 score, highlighting the notion that simple devices can be more accurately modeled.

**Table 4: NSR and F1 Scores**


| Data Set | Average NSR Complexity | Average F1 Score |
|----------|------------------------|------------------|
| Home     | 8.548                  | 0.901            |
| Lab      | 9.596                  | 0.879            |
| UNSW     | 7.170                  | 0.942            |
| SCADA    | 2.791                  | 0.975            |

Figure 8 shows the average F1 score for each averaged across the malware traffic. The upper left of Figure 8 shows the model using the non-normalized data. It is shown to illustrate the negative correlation between a device's NSR complexity and the F1 score. This supports the idea that simple devices can be more accurately modeled. In the upper right of the figure we see how the model is improved by normalizing the data. Normalizing generally improves the efficiency and accuracy of clustering algorithms, this is especially true when the clustering algorithm uses the squared Euclidean distance metric as is done in the DBSCAN algorithm.



**Figure 8: F1 Score vs NSR Complexity [Home]**

Figure 9 shows the average F1 score of each device against the seven types of malware traffic. Again, the upper left shows the non-normalized scores and the upper right scores based on normalized data.







**Figure 9: F1 Score vs NSR Complexity [Lab]**

Figure 10 and Figure 11 show the results for devices on the UNSW data set and SCADA dataset respectively. The UNSW results are consistent with the Home and Lab data sets in showing a negative correlation of F1 score and NSR score for non-normalized data.

The SCADA data set is distinctly different from the other three datasets, in that there does not appear to be an obvious correlation with complexity and modelability. The SCADA dataset had the lowest average NSR and the highest average F1 score. This is consistent with the notion that SCADA devices are very simple in terms of their network footprint and the results support this.



**Figure 10: F1 Score vs NSR Complexity [UNSW]**



**Figure 11: F1 Score vs NSR Complexity [SCADA]**

## 5. Conclusion

This research developed a method for measuring the complexity of IoT devices based on their network traffic. This method called the Noise to Signal Ratio (NSR) uses a clustering algorithm to determine how much of the traffic from a device can be classified as a signal and how much as noise. The number of clusters from this algorithm feeds a Gaussian mixture model that is used to construct a behavioral model for each device and classify normal versus abnormal traffic.

This model was then run against seven very different types of actual malware traffic to determine the efficacy of the model in classifying a device's normal traffic from malware traffic. The results show that the model is effective, with many devices having perfect F1 scores. The results also show that F1 scores are generally higher for less complex devices, supporting the claim that simple devices can be more accurately modeled than complex devices.

This suggests that automatic blocking of malware traffic could be done, especially on devices that are simple, i.e. have low NSR scores.

Future work will examine how the NSR of devices can be applied to other anomaly detection methods such as isolation forest and single class support vector machines.

|       |                                          |
|-------|------------------------------------------|
| A.I.  | Artificial intelligence                  |
| C&C   | command and control                      |
| DDOS  | distributed denial of service            |
| GMM   | Gaussian mixture method                  |
| IoT   | Internet of things                       |
| LSTM  | long term- short term                    |
| M.L.  | machine learning                         |
| NSR   | noise to signal ratio                    |
| NTC   | network traffic classifier               |
| RNN   | recurrent neural network                 |
| SCADA | Supervisory Control and Data Acquisition |

## Bibliography & References

1. **OpenConnectivity.** OCF Solving the IoT Standards Gap. [Online] 2020. <https://openconnectivity.org/>.
2. **Fagan, M, et al.** *Foundational Cybersecurity Activities for IoT Device Manufacturers*. s.l. : National Institute of Standards and Technology, 2020.
3. **Various.** *The C2 Consensus on IoT Device Security Baseline Capabilities*. s.l. : The Consumer Technology Association, 2019.
4. —. *CYBER; Cyber Security for Consumer Internet of Things*. s.l. : European Telecommunications Standards Institute, 2019.
5. **SB-327.** SB-327 Information privacy: connected devices. [Online] September 28, 2018. [https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill\\_id=201720180SB327](https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201720180SB327).

6. **Ripple20.** Overview- Ripple20. [Online] 2020. <https://www.jsf-tech.com/ripple20/>.
7. **Micro, New Mirai Variant - Trend.** New Mirai Variant Expands Arsenal, Exploits CVE-2020-10173. [Online] July 10, 2020. <https://blog.trendmicro.com/trendlabs-security-intelligence/new-mirai-variant-expands-arsenal-exploits-cve-2020-10173/>.
8. **Bezawada, Bruhadeshwar and Bachani, Maalvika and Peterson, Jordan and Shirazi, Hossein and Ray, Indrakshi and Ray, Indrajit.** Behavioral Fingerprinting of IoT Devices. *Proceedings of the 2018 Workshop on Attacks and Solutions in Hardware Security*. 2018, pp. 41-50.
9. **Lopez-Martin, Manuel, et al.** Network traffic classifier with convolutional and recurrent neural networks for Internet of Things. *IEEE Access*. 2017, pp. 18042-18050.
10. **Miettinen, Markus, et al.** IoT Sentinel: Automated device-type identification for security enforcement in IoT. *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. 2017, pp. 2177-2184.
11. **Marchal, Samuel, et al.** AuDI: Toward Autonomous IoT Device-Type Identification Using Periodic Communication. *IEEE Journal on Selected Areas in Communications*. 2019, pp. 1402-1412.
12. **Nguyen, Thien Duc, et al.** DIoT: A Federated Self-Learning Anomaly Detection System for IoT. *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. 2019, pp. 756-767.
13. **Ortiz, Jorge , Crawford, Catherine , Le, Franck.** DeviceMien: Network Device Behavior Modeling for Identifying Unknown IoT Devices. *Proceedings of the International Conference on Internet of Things Design and Implementation*. 2019, pp. 106-117.
14. **Parmisano, Agustin, Garcia, Sebastian and Erquiaga, Maria Jose.** Stratosphere Laboratory, A labeled dataset with malicious and benign IoT network traffic. [Online] April 20, 2020. <https://www.stratosphereips.org/datasets-iot23>.
15. **Various.** JOY: Network Capture and Analysis Tool. 2020.
16. **Ling, Robert F.** On the theory and construction of k-clusters. *The computer journal*. 1972, pp. 326-332.
17. **Akita.** SOHO & Home Cyber Security as a Service. *akita.cloud*. [Online] 2020. <https://shop.akita.cloud/>.
18. **Sentry, Cujo.** Network Security and Device Protection. [Online] 2020. <https://cujo.com/sentry/>.
19. **Bitdefender.** Smart Home Cyber Security for Your Business. [Online] 2020. <https://www.bitdefender.com/iot/>.



# **How DOCSIS Time Protocol makes the SYNC Specification Tick**

## **Automating the DTP Algorithm**

A Technical Paper prepared for SCTE•ISBE by

**Elías Chavarría Reyes, Ph.D.**

Senior Software Engineer  
Cisco Systems, Inc.  
San Jose, CA 95134  
408-527-7793  
elchavar@cisco.com

**John T. Chapman**

CTO Cable Access, Fellow  
Cisco Systems, Inc.  
San Jose, CA 95124  
408-526-7651  
jchapman@cisco.com

# Table of Contents

| Title                                                                | Page Number |
|----------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                 | 3           |
| 2. Market Opportunity.....                                           | 3           |
| 3. Why is time synchronization needed?.....                          | 4           |
| 4. DOCSIS Time Protocol (DTP).....                                   | 7           |
| 4.1. How timing protocols work .....                                 | 7           |
| 4.1. What DTP does .....                                             | 7           |
| 4.2. DTP Messaging.....                                              | 8           |
| 4.3. DTP – Determining the path delay .....                          | 9           |
| 5. DTP Timing Model.....                                             | 10          |
| 6. DTP Calibration .....                                             | 12          |
| 6.1. Measuring the CMTS-CM combined timing parameters .....          | 13          |
| 6.2. Measuring the HFC timing parameters .....                       | 14          |
| 6.3. Computing the DTP Time Adjustment in a live DOCSIS system ..... | 16          |
| 7. Requirements, Limitations and other aspects.....                  | 17          |
| 7.1. Requirements .....                                              | 17          |
| 7.2. Limitations .....                                               | 17          |
| 7.3. Other aspects .....                                             | 18          |
| 8. What's next?.....                                                 | 18          |
| 9. Conclusions.....                                                  | 19          |
| Abbreviations .....                                                  | 19          |
| Bibliography & References.....                                       | 21          |

## List of Figures

| Title                                                          | Page Number |
|----------------------------------------------------------------|-------------|
| Figure 1 – Heterogeneous Network .....                         | 3           |
| Figure 2 – Frequency vs Phase vs Time synchronization .....    | 4           |
| Figure 3 – Interference at UE and eNB.....                     | 5           |
| Figure 4 – Phase Synchronization in a Mobile Network .....     | 6           |
| Figure 5 – Timing Delivery over DOCSIS Backhaul .....          | 6           |
| Figure 6 – Two-way Time Transfer .....                         | 7           |
| Figure 7 – DOCSIS Time Protocol – System Overview .....        | 8           |
| Figure 8 – CMTS is DTP Master .....                            | 9           |
| Figure 9 – DTP Operation .....                                 | 9           |
| Figure 10 – DTP Timing Model .....                             | 10          |
| Figure 11 – Measuring CMTS-CM combined timing parameters ..... | 13          |
| Figure 12 – Measuring HFC timing parameters.....               | 15          |

## List of Tables

| Title                                                                    | Page Number |
|--------------------------------------------------------------------------|-------------|
| Table 1 – 4G/LTE and 5G Air Interface Synchronization Requirements ..... | 5           |
| Table 2 – DTP Delays .....                                               | 11          |

## 1. Introduction

CableLabs® released the I01 SYNC specification on April 20, 2020. The SYNC specification describes how to build and deploy an IEEE 1588/PTP participant DOCSIS network. The SYNC specification is targeted at the new and evolving mobile backhaul market.


The core technology that makes timing work is the DOCSIS Time Protocol (DTP). DTP was introduced in DOCSIS 3.1 and described in [1] but has not been deployed as of yet. The operation of DTP is not well understood and there are many new considerations in order to bring it from a specification to a deployed product. As a result of this strong interest by operators, there is a need to re-explore in detail several aspects of DTP, including:

- what DTP does,
- how it works,
- how it is calibrated,
- how it is used in a live system,
- how it interplays with other DOCSIS procedures,
- limitations, and tradeoffs.

The goal of this paper is to provide this detailed exploration of DTP and enable the Cable community to have a common understanding of DTP so they can confidently develop it and use it.

## 2. Market Opportunity

As highlighted in [2], mobile network operators (MNOs) are relying on densifying their radio access network (RAN) to improve the coverage, capacity, and throughput for both their 4G and, more importantly, 5G systems. As depicted in Figure 1, MNOs achieve this densification by installing low-powered base stations (called small cells) underneath their existing high-powered base stations (called macrocells), creating a heterogeneous network (HetNet).



**Figure 1 – Heterogeneous Network**

For small cells to work effectively, three key elements are required:

1. Location
2. Power
3. Backhaul

Cable operators are in an advantageous position to offering a solution for these three requirements, as discussed in detail in [2]. As such, the authors strongly believe that cable operators of today are the mobile network operators of tomorrow.

In terms of backhaul, these are the requirements:


1. high throughput to sustain the traffic of the small cells,
2. low latency and jitter, and
3. network timing.

In [2], the authors focused on how to achieve low latency backhaul over DOCSIS. In this paper, we focus on how to provide network timing by using the DOCSIS Time Protocol.

### 3. Why is time synchronization needed?

In this section we explore in detail the need for time and frequency synchronization among the elements of the RAN in a mobile network.


First, we will clarify the concepts of frequency, phase, and time synchronization. In Figure 2, we depict the difference among the three. Frequency synchronization refers to two elements that are beating at the same pace, but are not necessarily aligned at the instant when they beat. Phase synchronization refers to two elements that are beating at the same pace, and are aligned at the instant when they beat; however, the time they report at each beat is not necessarily the same. Time synchronization refers to two elements that are beating at the same pace, are aligned at the instant when they beat, and the time they report at each beat is the same. Throughout this paper we will use the term “timing synchronization” to refer to frequency, phase, and time synchronization all together unless explicitly stated.



**Figure 2 – Frequency vs Phase vs Time synchronization**

Many mobile networks previously required only frequency synchronization. However, 4G and 5G networks require also phase and time synchronization for several reasons, including:

- a) Use of Time Division Duplexing (TDD): In TDD mode, the base stations (also called eNBs) utilize a same frequency for both upstream and downstream transmission. A single base station will switch in time between transmitting and receiving data. Since base stations have overlapping coverage, particularly in a HetNet, those base stations need to be aligned in terms of when to be in transmission mode, or in reception mode; otherwise, they will cause interference to each other, as depicted in Figure 3. In Figure 3, the base stations are completely misaligned, causing interference between user equipment (UE) at the cell-edge, and also at the base stations themselves.



**Figure 3 – Interference at UE and eNB**


- b) Use of advanced base station coordination techniques: In both 4G and 5G systems, new techniques were introduced to either mitigate potential interference among base stations operating in the same frequency band (e.g., enhanced intercell interference cancellation, or eICIC), or improve the throughput achievable by UEs under the coverage of multiple base stations (e.g., Coordinated Multi Point operation, or CoMP). For these techniques to work effectively, the base stations need to be tightly synchronized.

Table 1 [3] summarizes the synchronization requirements for 4G LTE Frequency-Division Duplex (FDD), 4G LTE TDD, and 5G TDD.

**Table 1 – 4G/LTE and 5G Air Interface Synchronization Requirements**

| Table Heading | Frequency                                                                      | Phase                                                                                                                                                                                                         |                                                                          |
|---------------|--------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| 4G LTE FDD    | $\pm 50$ ppb                                                                   | None                                                                                                                                                                                                          | 3GPP TS 36.104 §6.5.1                                                    |
| 4G LTE TDD    | $\pm 50$ ppb (wide area)<br>$\pm 100$ ppb (local area)<br>$\pm 250$ ppb (home) | $10 \mu\text{s}$ (wide: cell radius $>3$ km)<br>$3 \mu\text{s}$ (local: cell radius $<3$ km)<br>$1.33 \mu\text{s} + T_{\text{prop}}$ (home eNB radius $>500$ m)<br>$3 \mu\text{s}$ (home eNB radius $<500$ m) | Phase:<br>3GPP TS 36.133 §7.4.2<br>Frequency:<br>3GPP TS 36.922 §6.4.1.2 |
| 5G TDD        | $\pm 50$ ppb (wide area)<br>$\pm 100$ ppb (local area)                         | $\leq 3 \mu\text{s}$                                                                                                                                                                                          | 3GPP TS 38.104 Table 6.5.1.2.1                                           |

The SYNC specification was designed with the goal of achieving the  $3 \mu\text{s}$  phase requirement. The key concept to note is that the  $\pm 3 \mu\text{s}$  requirement refers to the allowed time error between the RAN network elements. In a deployed system, this  $3 \mu\text{s}$  requirement is typically achieved by setting a  $\pm 1.5 \mu\text{s}$  phase requirement between the RAN network elements and a common time source, as depicted in Figure 4.




**Figure 4 – Phase Synchronization in a Mobile Network**

One way to achieve the timing synchronization is to use a Global Navigation Satellite System (GNSS) as time source, and equip every base station with a GNSS receiver. This works well for macrocells. Macrocells are typically deployed in outdoor environments, and the cost of the GNSS receiver is very small compared to the cost of the macrocell itself. For small cells, using GNSS may not work well for two reasons: GNSS signal availability and cost of GNSS receiver.

Small cells are meant to be low-cost devices deployed in large quantities and may be deployed in indoor environments where the signal from a GNSS could be unreliable. Even where the GNSS signal is reliable, the cost of the GNSS receiver might significantly increase the cost of the small cell. Because of these two reasons, network-based timing protocols, which do not depend on GNSS signal availability and have lower cost impact on small cells are preferable.

What Figure 4 shows with a single arrow from the base stations to the time source is actually comprised of the DOCSIS backhaul and other network elements, as depicted in Figure 5.



**Figure 5 – Timing Delivery over DOCSIS Backhaul**

Figure 5 shows a grandmaster (GM) synchronized to a GNSS. Through a series of network elements (marked as E), the timing information is delivered to the CMTS through the use of 1588/PTP and, optionally, SyncE. The CMTS utilizes existing DOCSIS timing properties and the DOCSIS Time Protocol (DTP) to convey the information that the CM needs to be properly synchronized. The CM then delivers the timing information to the base station, called eNB, using 1588/PTP and, optionally, SyncE. The end application resides within the eNB.

Comparing Figure 4 and Figure 5, we see that the DOCSIS backhaul, i.e., the CMTS and CM, is only one of the elements between the time source, i.e., GNSS, and the end application. Thus, the 1.5  $\mu$ s budget is actually distributed among the DOCSIS elements and the rest of the elements sitting between the GNSS and the end application.

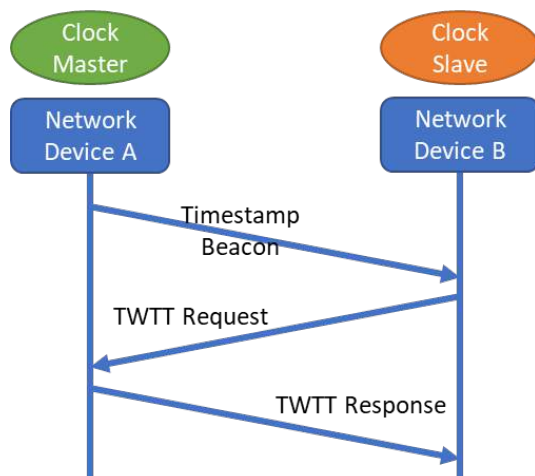
In the next section, we discuss the details of how DTP works.

## 4. DOCSIS Time Protocol (DTP)

### 4.1. How timing protocols work

Before diving into the specifics of DTP, we will provide a high-level view of how timing protocols work.

All timing protocols use a variant of a Two-Way Time Transfer (TWTT), as depicted in Figure 6.




**Figure 6 – Two-way Time Transfer**

There are two devices, one containing the master clock (Network Device A), and one containing the slave clock (Network Device B). Network Device A sends a timestamp to Network Device B. This timestamp will arrive at Network Device B at some time after it was created by Network Device A. In other words, the timestamp is delayed. Network Device A and B exchange messages to determine the network delay and asymmetry. Once these two factors are identified, Network Device B corrects its timestamp to match the one of Network Device A. Even though this procedure was used to synchronize the time, the same mechanism can be used to synchronize the frequency.

### 4.1. What DTP does

In its simplest form, DTP allows the existing timing and frequency system of DOCSIS to be interfaced to external timing protocols with high accuracy, as shown in Figure 7. Once the CMTS has a frequency and time source synchronized to an external source, DTP allows the source to be replicated at the egress port of the CM.



**Figure 7 – DOCSIS Time Protocol – System Overview**

As highlighted in [4], for the DOCSIS DTP frequency path, the CMTS PLL (Phase-Locked Loop) locks onto the frequency component of the external timing protocol, e.g., SyncE or IEEE-1588. The output of the CMTS PLL drives the DOCSIS downstream baud rate. The CM receives the baud rate frequency and locks to it with its PLL. The CM PLL then drives the frequency output on the CMCI port with SyncE and/or IEEE-1588.

As highlighted in [4], for the DOCSIS DTP timestamp path, the CMTS synchronizes its DOCSIS 3.1 Extended Timestamp to the timestamp of the external protocol, e.g., IEEE-1588. The DOCSIS Extended Timestamp is sent to the CM as part of the DOCSIS protocol where it is converted back to any desired format, e.g., IEEE-1588. The DTP protocol that runs between the CMTS and the CM computes the time delay in the downstream path while taking into account the asymmetry of the DOCSIS system. This delay is then added to the timestamp in the CM so that the timestamp that is sent out from the CM closely matches the timestamp received by the CMTS.

The main takeaways from the above two paragraphs are:

- The CMTS and CM are already synchronized in frequency. The CMTS frequency will need to be synchronized to an external frequency source. The CM will need to pass its local frequency to its CMCI port.
- The CMTS and CM are already phase aligned, but not time synchronized. DTP provides the means to achieving the time synchronization.

## 4.2. DTP Messaging

Similarly to how other time protocols operate (see Section 4.1), DTP operates under a master/slave approach. The DTP master is in charge of:


- a) Initiating DTP message exchange, and
- b) Computing the time adjustment for the CM.

The MULPI spec [4] allows for either the CMTS or the CM to be the DTP master. Figure 8 depicts the message flow between the DTP master and DTP slave for the case of the CMTS being the DTP master. The DTP request (DTP-REQ) provides configuration information to the CM regarding the CMTS and HFC timing values. The CM does not need these values to compute the ones that it sends back in the DTP response (DTP-RSP).

The key value reported back by the CM is the true ranging offset (TRO) and the CM timing values. The CMTS uses this data, the CMTS timing values, and HFC timing values to compute the DTP time



adjustment, and then sends this time adjustment to the CM in a DTP info (DTP-INFO) message. Once the CM receives the DTP-INFO, the CM replies back with a DTP acknowledgement (DTP-ACK).

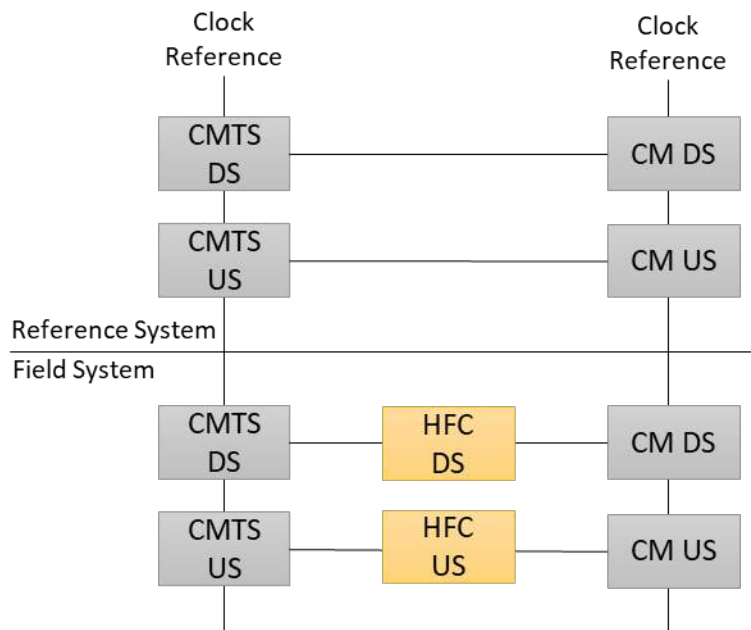


**Figure 8 – CMTS is DTP Master**

The time interval at which the DTP message exchange is repeated is configurable, with 10 seconds being the minimum value and no maximum value.

### 4.3. DTP – Determining the path delay

At a high-level, DTP relies on the exchange of messages (see Section 4.2) to determine the forward path delay needed to compensate the time at the CM.



**Figure 9 – DTP Operation**

In Figure 9, we depict a simplified field system composed of a CMTS (downstream – DS, and upstream – US), and HFC plant (DS and US), and CM (DS and US). In this figure, the forward path delay is composed by the CMTS DS, HFC DS, and CM DS elements. The underlying assumptions used to compute these elements are:

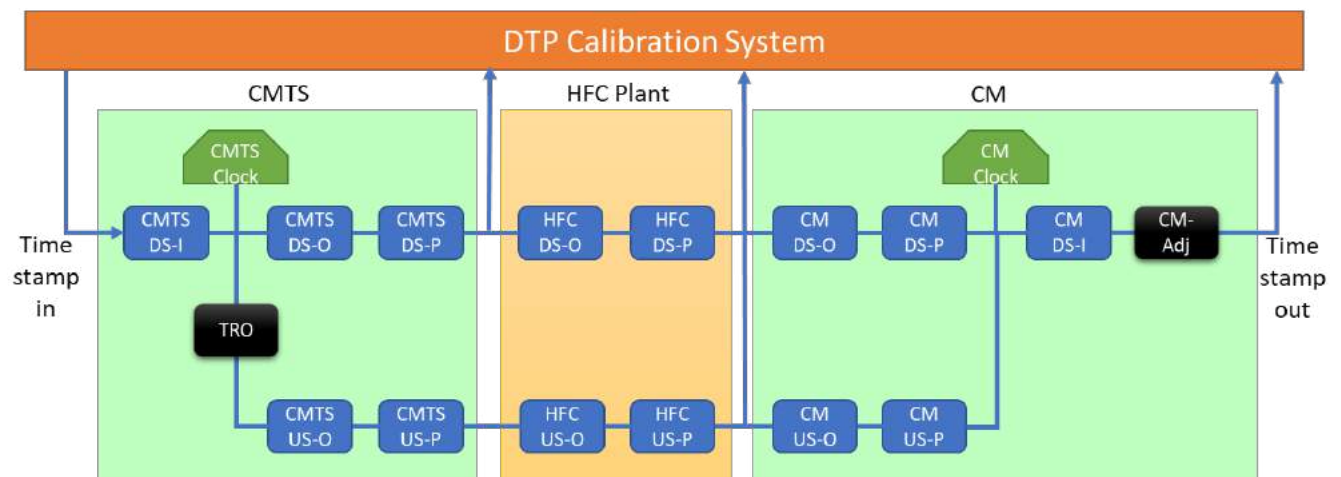
1. The CMTS and CM parameters can be measured offline.
2. The HFC path delay in downstream is equal to the HFC upstream path delay. Note that this assumption is not always true; however, if the error is within the error budget allocated to the HFC, then the error can be ignored.
3. The difference in path delay between a field system and a zero-length plant is attributable to the total HFC path delay (DS & US).

With these assumptions in place, it follows that by computing the difference in path delay between a field system and a zero-length plant, we can obtain the total HFC path delay. We take this value and divide it by two, and we get the one-way HFC path delay. Now, because the CMTS DS and CM DS parameters were calibrated offline, the overall path delay can now be computed as the sum of CMTS DS, HFC DS path delay, and CM DS.

So, for the above to work, we need to have the CMTS and CM parameters calibrated offline, which an entity such as CableLabs® could lead. In the next section, we dive deeper into what exactly the CMTS, CM, and HFC parameters represent.

## 5. DTP Timing Model

The DTP timing model is depicted in Figure 10.



**Figure 10 – DTP Timing Model**

The timing values have blue background. The ones associated with the CMTS are preceded by the word “CMTS”, the ones associated with the HFC are preceded by the word “HFC”, and the ones associated with the CM are preceded by the word “CM”. The parameters associated with downstream include the word “DS”, and the ones associated with upstream include the word “US”. There are three types of parameters:

1. Interface Delay: Include the suffix “-I”
2. Path: Include the suffix “-P”. This is a vendor specific characterized delay of the physical circuit path. For the CMTS and CM, this may be a measured or calibrated value that is supplied as part of calibration. For the HFC plant, this is the value that is calculated as part of the DTP calculations, as discussed in Section Figure 9.
3. Offset: Include the suffix “-O”. This is a known offset due to interleaving or some other configuration.

Table 2 summarizes all the parameters, as defined in [4]. Compared to Figure 10, the spec uses the additional prefix “t-” to refer to the timing parameters.

**Table 2 – DTP Delays**

|      | <b>Delay Type</b> | <b>Delay Name</b> | <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------|-------------------|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CMTS | Interface         | t-cmts-ds-i       | This is the circuit delay from the CMTS clock input interface (DTI or NSI) to the internal CMTS timestamp reference point. This is a manufacturer's value and is supplied by the CMTS.                                                                                                                                                                                                                                                                                                        |
|      | Path              | t-cmts-ds-p       | This is the intrinsic path delay contribution from the CMTS timestamp reference point to the CMTS downstream PHY output. This is a measured value and supplied by the CMTS.                                                                                                                                                                                                                                                                                                                   |
|      |                   | t-cmts-us-p       | This is the intrinsic path delay contribution from the CMTS PHY upstream input to the CMTS timestamp reference point. This is a measured value and supplied by the CMTS.                                                                                                                                                                                                                                                                                                                      |
|      | Offset            | t-cmts-ds-o       | This is the known delay contribution in the downstream CMTS PHY path that is associated with configuration elements such as interleaving. This value is known and is supplied by the CMTS.                                                                                                                                                                                                                                                                                                    |
|      |                   | t-cmts-us-o       | This is the known delay contribution in the downstream CMTS PHY path that is associated with configuration elements such as interleaving. This value is known and is supplied by the CMTS.                                                                                                                                                                                                                                                                                                    |
| HFC  | Path              | t-hfc-ds-p        | This is the intrinsic path delay of the fiber and coax elements of the HFC plant. The DTP algorithm calculates this value.                                                                                                                                                                                                                                                                                                                                                                    |
|      |                   | t-hfc-us-p        | This is the intrinsic path delay of the fiber and coax elements of the HFC plant exclusive of fixed delay elements. The DTP algorithms calculate this value. The basic DTP algorithm assumes the upstream and downstream path delay are equal by using the offset values to compensate for fixed and asymmetrical delays.                                                                                                                                                                     |
|      | Offset            | t-hfc-ds-o        | This delay represents any fixed delay elements in the HFC path that contribute to delay. One example may be a digitization circuit, optical node and amplifier circuit delays. This value may be unique per HFC path due to different path elements. This value is supplied by the CMTS.<br>By specifying appropriate HFC downstream and upstream offset values correctly and by setting the asymmetry appropriately, the HFC downstream and upstream path delays can be assumed to be equal. |
|      |                   | t-hfc-us-o        | This delay represents any fixed delay elements in the HFC path that contribute to delay. One example may be a digitization circuit, optical node and amplifier circuit delays. This value may be unique per HFC path due to different path elements. This value is supplied by the CMTS.                                                                                                                                                                                                      |
| CM   | Interface         | t-cm-ds-i         | This is the circuit delay from the internal CM timestamp reference point to the clock output interface (CMCI). This value is manufacturer's value and is supplied by the CM or by a CMTS override.                                                                                                                                                                                                                                                                                            |
|      | Path              | t-cm-ds-p         | This is the intrinsic path delay contribution from the CM PHY downstream input to the CM timestamp reference point. This is a measured value and supplied by the CM or by a CMTS override.                                                                                                                                                                                                                                                                                                    |
|      |                   | t-cm-us-p         | This is the intrinsic path delay contribution from the CM timestamp reference point to the CM PHY upstream output. This is a measured value and supplied by the CM or by a CMTS override.                                                                                                                                                                                                                                                                                                     |
|      | Offset            | t-cm-ds-o         | This is the known delay contribution in the downstream CM PHY path that is associated with configuration elements such as interleaving. This value is known and is supplied by the CM or by a CMTS override.                                                                                                                                                                                                                                                                                  |

|  | Delay Type | Delay Name | Description                                                                                                                                                                                                |
|--|------------|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  |            | t-cm-us-o  | This is the known delay contribution in the upstream CM PHY path that is associated with configuration elements such as interleaving. This value is known and is supplied by the CM or by a CMTS override. |
|  | Other      | t-cm-adj   | This is the value that must be added to the CM unadjusted timestamp to have the CM timestamp be equal to the CMTS timestamp in real time. This value is calculated by the DTP Master.                      |

In the addition to the timing values described above, the other parameter of relevance is the True Ranging Offset (TRO). The TRO is the measured ranging offset of the CM between two defined reference points. TRO is a measured (or derived) value that is different than the actual implemented ranging offset a CM might use in its communication with the CMTS. A key property of the TRO is that the value of TRO is the equivalent to the round-trip delay of the combined downstream and upstream propagation delays of the HFC plant, the CMTS and CM PHY paths. For more details on the TRO, see [4].

With the TRO, CMTS timing values, CM timing values, and HFC offset values, the DTP math described in MULPI [4] provides the formulas to compute the necessary time adjustment for the CM. Those formulas are shown here for reference:

- (1)  $t-tro = t-cmts-ds-o + t-cmts-ds-p + t-hfc-ds-o + t-hfc-ds-p + t-cm-ds-o + t-cm-ds-p + t-cm-us-o + t-cm-us-p + t-hfc-us-o + t-hfc-us-p + t-cmts-us-o + t-cmts-us-p$
- (2)  $t-cm-adj = t-cmts-ds-i + t-cmts-ds-o + t-cmts-ds-p + t-hfc-ds-o + t-hfc-ds-p + t-cm-ds-o + t-cm-ds-p + t-cm-ds-i$

The limitation is that there does not exist any measurement equipment that, as of the time of the writing of this paper, can measure separately the CMTS, CM, and HFC timing values. Until the testing limitation is addressed by test equipment vendors, rather than characterizing the CMTS and CM separately, they must be characterized jointly for specific configurations, as discussed in the next section.

## 6. DTP Calibration

The approach for DTP calibration is based on the following assumptions:

1. The timing values of the CM and CMTS will be characterized jointly due to lack of testing equipment, as of the time of the writing of this paper, that can characterize them separately.
2. The joint characterization of the CM and CMTS needs to be done for every configuration parameter, e.g., interleaving, that affects the upstream or downstream delay.
3. The timing values of an HFC element can be characterized indirectly by comparing timing behavior of a DOCSIS network without such an HFC element and with such an HFC element.


With the above assumptions, the SYNC spec provides the following three steps to perform the DTP calibration and using the measured values on a live system.

1. Measuring the CMTS-CM combined timing parameters: This is done at a reference-length plant. Ideally the plant would be zero-length, but it is sufficient for the plant to be as simple as possible and with a known length.

2. Measuring the HFC timing parameters: This is done indirectly by inserting the HFC element between the CMTS and CM, and computing the difference between this plant and the reference-length plant.
3. Computing the DTP Time Adjustment in a live DOCSIS system. This is done based on the values computed or measured in the previous two steps.

### 6.1. Measuring the CMTS-CM combined timing parameters

The goal is to jointly characterize the CMTS-CM timing parameters using a reference-length plant, i.e., the HFC path is just a coax connection of reference length from the CMTS to the CM, as depicted in Figure 11.



**Figure 11 – Measuring CMTS-CM combined timing parameters**

In the following formulas, we utilize the suffix “-R” to refer to any values that are specific to the reference-length plant.

For a reference-length plant, formulas (1) and (2) become:

$$\begin{aligned}
 (3) \quad t_{tro-R} &= t_{cmts-ds-o} + t_{cmts-ds-p} + t_{hfc-ds-o-R} + t_{hfc-ds-p-R} + t_{cm-ds-o} + t_{cm-ds-p} + t_{cm-us-o} + t_{cm-us-p} \\
 &\quad + t_{hfc-us-o-R} + t_{hfc-us-p-R} + t_{cmts-us-o} + t_{cmts-us-p} \\
 (4) \quad t_{cm-adj-R} &= t_{cmts-ds-i} + t_{cmts-ds-o} + t_{cmts-ds-p} + t_{hfc-ds-o-R} + t_{hfc-ds-p-R} + t_{cm-ds-o} + t_{cm-ds-p} \\
 &\quad + t_{cm-ds-i}
 \end{aligned}$$

Note that the CMTS and CM parameters do not require a “-R” suffix because their values are independent of the length of the HFC in the reference-length plant.

Since the HFC is just a coax connection in the reference-length plant, it is assumed that the HFC upstream and downstream offset timing values are both zero, i.e.,  $t_{hfc-us-o-R} = 0$  and  $t_{hfc-ds-o-R} = 0$ . Thus, formulas (3) and (4) become:

$$\begin{aligned}
 (5) \quad t_{tro-R} &= t_{cmts-ds-o} + t_{cmts-ds-p} + t_{hfc-ds-p-R} + t_{cm-ds-o} + t_{cm-ds-p} + t_{cm-us-o} + t_{cm-us-p} + t_{hfc-us-p-R} \\
 &\quad + t_{cmts-us-o} + t_{cmts-us-p} \\
 (6) \quad t_{cm-adj-R} &= t_{cmts-ds-i} + t_{cmts-ds-o} + t_{cmts-ds-p} + t_{hfc-ds-p-R} + t_{cm-ds-o} + t_{cm-ds-p} + t_{cm-ds-i}
 \end{aligned}$$

Now, we record the following values for the reference-length plant:

- $t_{tro-R}$ : True ranging offset as reported by the CM.

- *t-cm-adj-R*: Value of the DTP time adjustment that brings the average PTP 2-Way Time Error (cTE) to zero.
- *t-hfc-ds-p-R*: Downstream path delay introduced by the coax cable going from the CMTS to the CM. This is computed based on the length of the coax cable and the speed of propagation in the coax cable. Typically, this speed is around 1.5ns/ft.
- *t-hfc-us-p-R*: Downstream path delay introduced by the coax cable going from the CMTS to the CM. This is computed based on the length of the coax cable and the speed of propagation in the coax cable. Typically, this speed is around 1.5ns/ft.

Rearranging the terms in formulas (5) and (6), we get

$$(7) \quad t-tro-R - t-hfc-ds-p-R - t-hfc-us-p-R = t-cmts-ds-o + t-cmts-ds-p + t-cmts-us-o + t-cmts-us-p + t-cm-ds-o + t-cm-ds-p + t-cm-us-o + t-cm-us-p$$

$$(8) \quad t-cm-adj-R - t-hfc-ds-p-R = t-cmts-ds-i + t-cmts-ds-o + t-cmts-ds-p + t-cm-ds-o + t-cm-ds-p + t-cm-ds-i$$

The most important aspect of the formulas (7) and (8) is that on the left-hand side of each formula we have values that are either measured or calculated based on the reference-length plant, and on the right-hand side we have a combination of CMTS and CM timing parameters. Thus, we achieved the goal of measuring the CMTS-CM combined timing parameters.

Now that we have an expression that summarizes the combination of CMTS and CM parameters (right-hand side of formulas (7) and (8)), we use formulas (7) and (8) and combine those with the generic formulas (1) and (2). Here are the formulas (1) and (2) once more, with the CMTS and CM parameters moved as to group them for simplicity:

$$(1) \quad t-tro = t-hfc-ds-o + t-hfc-ds-p + t-hfc-us-o + t-hfc-us-p + (t-cmts-ds-o + t-cmts-ds-p + t-cmts-us-o + t-cmts-us-p + t-cm-ds-o + t-cm-ds-p + t-cm-us-o + t-cm-us-p)$$

$$(2) \quad t-cm-adj = t-hfc-ds-o + t-hfc-ds-p + (t-cmts-ds-i + t-cmts-ds-o + t-cmts-ds-p + t-cm-ds-o + t-cm-ds-p + t-cm-ds-i)$$

The terms in bold in formula (1) are the same as the right-hand side of formula (7). The terms in bold in formula (2) are the same as the right-hand side of formula (8). So, plugging formulas (7) and (8) into formulas (1) and formula (2), we get new generic formulas:


$$(9) \quad t-tro = t-hfc-ds-o + t-hfc-ds-p + t-hfc-us-o + t-hfc-us-p + (t-tro-R - t-hfc-ds-p-R - t-hfc-us-p-R)$$

$$(10) \quad t-cm-adj = t-hfc-ds-o + t-hfc-ds-p + (t-cm-adj-R - t-hfc-ds-p-R)$$

## 6.2. Measuring the HFC timing parameters

The goal is to characterize the timing parameters of a single HFC element. While it is possible to characterize a set of HFC elements as a single unit, the results would only be usable when that same set is present in a deployment.

First, the HFC element of interest is inserted between the CMTS and the CM, as depicted in Figure 12. For the plant described in Figure 12, we utilize the suffix “-1” to refer to any values that are specific to this plant. This represents the plant under test in an actual deployment. The goal is to generate a formula for CM-adj.



**Figure 12 – Measuring HFC timing parameters**

In Figure 12, the HFC is composed of three parts. Part 1 is the cable going from the CMTS to the HFC element. Part 2 is the HFC element itself. Part 3 is the cable going from the HFC element to the CM.

Now, we record the following values for the plant in Figure 12:

- $t\text{-tro-}l$ : True ranging offset as reported by the CM.
- $t\text{-cm-adj-}l$ : Value of the DTP time adjustment that brings the average PTP 2-Way Time Error (cTE) to zero
- $t\text{-hfc-ds-p-}l$ : Downstream path delay introduced by the fiber and coax cables going from the CMTS to the CM. This value includes the downstream path delay from the CMTS to the HFC element and from the HFC element to the CM. This value is computed based on the length of the fiber and coax cables and the speed of propagation in each.
- $t\text{-hfc-us-p-}l$ : Upstream path delay introduced by the fiber and coax cables going from the CMTS to the CM. This value includes the upstream path delay from the CM to the HFC element and from the HFC element to the CMTS. This is computed based on the length of the fiber and coax cables and the speed of propagation in each.

Note that there is no “-l” parameter for the HFC offset because the HFC offset parameter is independent of the coax and fiber lengths between the CMTS and the CM.

With the above values recorded, the following is obtained from formulas (9) and (10):

$$(11) \ t\text{-tro-}l = t\text{-hfc-ds-o} + t\text{-hfc-ds-p-}l + t\text{-hfc-us-o} + t\text{-hfc-us-p-}l + (t\text{-tro-R} - t\text{-hfc-ds-p-R} - t\text{-hfc-us-p-R})$$

$$(12) \ t\text{-cm-adj-}l = t\text{-hfc-ds-o} + t\text{-hfc-ds-p-}l + (t\text{-cm-adj-R} - t\text{-hfc-ds-p-R})$$

Rearranging the terms in formula (11), we get:

$$(13) \ t\text{-hfc-ds-o} + t\text{-hfc-us-o} = t\text{-tro-}l - t\text{-hfc-ds-p-}l - t\text{-hfc-us-p-}l - (t\text{-tro-R} - t\text{-hfc-ds-p-R} - t\text{-hfc-us-p-R})$$

Rearranging the terms in formula (12), we get:

$$(14) \ t\text{-hfc-ds-o} = t\text{-cm-adj-}l - t\text{-hfc-ds-p-}l - (t\text{-cm-adj-R} - t\text{-hfc-ds-p-R})$$

Plugging formula (14) into formula (13), we get

$$(15) t\text{-cm-adj-}l - t\text{-hfc-ds-p-}l - (t\text{-cm-adj-}R - t\text{-hfc-ds-p-}R) + t\text{-hfc-us-o} = t\text{-tro-}l - t\text{-hfc-ds-p-}l - t\text{-hfc-us-p-}l - (t\text{-tro-}R - t\text{-hfc-ds-p-}R - t\text{-hfc-us-p-}R)$$

Now, simplifying formula (15), we get:

$$(16) t\text{-hfc-us-o} = (t\text{-tro-}l - t\text{-hfc-us-p-}l) - (t\text{-cm-adj-}l - t\text{-cm-adj-}R) - (t\text{-tro-}R - t\text{-hfc-us-p-}R)$$

The most important aspect of formulas (14) and (16) is that on the left-hand side of each we have the HFC offset parameters, and on the right-hand side we have values that were either obtained from the reference-length plant (the “-R” values), or measured/calculated in the plant with the added HFC elements (the “-l” values). Thus, we get a numeric value for the HFC offset timing values.

The second most important aspect of the above results is that the values obtained for the HFC offset timing values can later be used with any other CMTS-CM pair different from the ones used for the reference-length plant.

Note that if there are multiple HFC elements, then the values of  $t\text{-hfc-ds-o}$  and  $t\text{-hfc-us-o}$  should reflect the addition of the timing parameters of each HFC element.

An operator wanting to use DTP should have all its HFC elements characterized from a DTP point of view. Not doing such characterization would eat away from the overall time error budget allocated to the DOCSIS network. Alternatively, a rough characterization might work as there may be enough margin in the time error budget, or the length of the fiber plant dominates the overall time error and negates small changes on the coax plant. So, the difference in calibration may be the difference between a good result and a great result. Is good, good enough? It depends on the result and the cost.

### 6.3. Computing the DTP Time Adjustment in a live DOCSIS system

Now that we have the CMTS-CM pair and HFC elements characterized, we can use those values to compute the DTP time adjustment in a live system.

Note: The variables  $t\text{-hfc-ds-o}$  and  $t\text{-hfc-us-o}$  are chosen to model both fixed delays and any path asymmetry between the upstream and downstream HFC transmission paths. This allows the assumption to be made that the remaining path delay from the hfc downstream path and the hfc upstream paths are equal. Hence,  $t\text{-hfc-us-p} = t\text{-hfc-ds-p}$ .

Updating formula (9) based on  $t\text{-hfc-us-p} = t\text{-hfc-ds-p}$ , we get:

$$(17) t\text{-tro} = t\text{-hfc-ds-o} + t\text{-hfc-ds-p} + t\text{-hfc-us-o} + t\text{-hfc-ds-p} + (t\text{-tro-}R - t\text{-hfc-ds-p-}R - t\text{-hfc-us-p-}R)$$

$$(18) t\text{-tro} = t\text{-hfc-ds-o} + 2*t\text{-hfc-ds-p} + t\text{-hfc-us-o} + (t\text{-tro-}R - t\text{-hfc-ds-p-}R - t\text{-hfc-us-p-}R)$$

Rearranging the terms in formula (18), we get  $t\text{-hfc-ds-p}$ :

$$(19) t\text{-hfc-ds-p} = [t\text{-tro} - t\text{-hfc-ds-o} - t\text{-hfc-us-o} - (t\text{-tro-}R - t\text{-hfc-ds-p-}R - t\text{-hfc-us-p-}R)]/2$$

Plugging formula (19) into formula (10), we get:

$$(20) t\text{-cm-adj} = t\text{-hfc-ds-o} + [t\text{-tro} - t\text{-hfc-ds-o} - t\text{-hfc-us-o} - (t\text{-tro-}R - t\text{-hfc-ds-p-}R - t\text{-hfc-us-p-}R)]/2 + (t\text{-cm-adj-}R - t\text{-hfc-ds-p-}R)$$



Distributing the denominator and signs, we get:

$$(21) t\text{-cm-adj} = t\text{-hfc-ds-o} + t\text{-tro}/2 - t\text{-hfc-ds-o}/2 - t\text{-hfc-us-o}/2 - t\text{-tro-R}/2 + t\text{-hfc-ds-p-R}/2 + t\text{-hfc-us-p-R}/2 + t\text{-cm-adj-R} - t\text{-hfc-ds-p-R}$$

Re-arranging and grouping like terms

$$(22) t\text{-cm-adj} = t\text{-cm-adj-R} + t\text{-tro}/2 - t\text{-tro-R}/2 + t\text{-hfc-ds-o} - t\text{-hfc-ds-o}/2 - t\text{-hfc-us-o}/2 - t\text{-hfc-ds-p-R} + t\text{-hfc-ds-p-R}/2 + t\text{-hfc-us-p-R}/2$$

Subtracting like terms

$$(23) t\text{-cm-adj} = t\text{-cm-adj-R} + t\text{-tro}/2 - t\text{-tro-R}/2 + t\text{-hfc-ds-o}/2 - t\text{-hfc-us-o}/2 - t\text{-hfc-ds-p-R}/2 + t\text{-hfc-us-p-R}/2$$

Further assuming that  $t\text{-hfc-ds-p-R} = t\text{-hfc-us-p-R}$ , we get

$$(24) t\text{-cm-adj} = t\text{-cm-adj-R} + t\text{-tro}/2 - t\text{-tro-R}/2 + t\text{-hfc-ds-o}/2 - t\text{-hfc-us-o}/2 - t\text{-hfc-us-p-R}/2 + t\text{-hfc-us-p-R}/2$$

$$(25) t\text{-cm-adj} = t\text{-cm-adj-R} + t\text{-tro}/2 - t\text{-tro-R}/2 + t\text{-hfc-ds-o}/2 - t\text{-hfc-us-o}/2$$

Now, pulling the denominator back out and group terms for usability, we end up with the final formula for  $t\text{-cm-adj}$

$$(26) t\text{-cm-adj} = t\text{-cm-adj-R} + [(t\text{-tro} - t\text{-tro-R}) + (t\text{-hfc-ds-o} - t\text{-hfc-us-o})]/2$$

The most important aspect is that the only element on the right-hand side that was not a pre-calibrated numeric value is the TRO ( $t\text{-tro}$ ) of the live system. The second important aspect that the formula hints at is that the DTP master, which is the entity computing the DTP time adjustment ( $t\text{-cm-adj}$ ) should know what HFC parameters to use. In other words, some entity, e.g., the DTP master, should have an association between the CM and the HFC elements between the such CM and the CMTS. With this association, then the right HFC timing values could be recovered.

## 7. Requirements, Limitations and other aspects

In this section, we capture requirements, limitations, and tradeoffs related to the use of DTP.

### 7.1. Requirements

DTP relies on the true ranging offset (TRO), which is itself related to the ranging offset. As such, for the proper operation of DTP, the CM must have properly ranged. Moreover, the DTP message exchange should be repeated at least every time that the ranging offset changes.

The accuracy of DTP relies on the accuracy of the DOCSIS extended timestamp, which is obtained by the CM from an OFDM primary downstream channel. Thus, for the proper operation of DTP, the system must have an OFDM primary downstream channel. From the CM side, not only D3.1 support is needed but also DTP support (since its support is optional), and PTP support on its CMCI interface.

### 7.2. Limitations

The focus of DTP is on providing the CM with the appropriate time adjustment to be synchronized to an external time reference. However, what protocols such as PTP also provide is traceability information of the timing source and the path to the timing source. DTP does not provide this traceability. Rather than

making DTP aware of PTP-specifics, the SYNC committee decided to let the PTP traceability information be sent from the CMTS (acting as a boundary clock) to the CM (acting as a boundary clock). DTP does not need to modify the traceability information, but DOCSIS does need to transport it..

As discussed in Section 6, the calibration of the DTP timing values must be done jointly for the CMTS and the CM, and must take into account any configuration values, such as interleaving, that affects those timing values. The output of this calibration will be two-fold: a list of timing values per CMTS-CM pair and per configuration, and a list of timing values per HFC element per configuration. The industry can choose one of these two approaches to carry out the calibration efforts:

- a) Let every MSO, CMTS, or CM vendor build its own calibration facilities and perform its own calibration measurements, or
- b) Have a central entity, such as CableLabs®, be the one in charge of setting the calibration facilities and performing the calibration measurements.

It is the opinion of the authors that the first approach will be initially used by the CM and CMTS vendors releasing DTP support. However, for scalability and consistency purposes, it is the opinion of the authors that the second approach is better.

The I01 version of the SYNC specification was released on April 20, 2020. It includes a robust description of the overall system and its operation when the external timing sources are using PTP profile G.8275.1. The support for G.8275.2 is work in progress and expected to be added in a future release; this works includes the allocation of time error budget for the network elements and the definition of the transport mechanism for the PTP Announce and Signaling messages.

While the SYNC specification is not completed yet for G.8275.2, there is nothing preventing DTP to be tested in a network using G.8275.2 for the purposes of measuring the overall time error; the caveat is that the traceability information would not be available at the CM. In other words, DTP itself is agnostic to the PTP profile being used by the CMTS or the CM themselves.

### **7.3. Other aspects**

The SYNC specification not only deals with the details of DTP, but also with the suggested allocation of the time error budget for the elements in the network. As discussed in Section 3, the DOCSIS network elements need part of the overall time error budget, but cannot be allocated all of it.

Each MSO needs to characterize the performance of their existing networks (from a time error point of view), including its non-DOCSIS elements, to validate if their existing networks can meet the total timing requirements or not. MSOs may find out that they need to upgrade parts of their networks, DOCSIS or non-DOCSIS elements, to meet the timing requirements.

For DTP to work effectively, the recovery of the DTP values by the DOCSIS system must be automated. As mentioned in Section 6.3, the DOCSIS system not only needs to recover the timing values of the CMTS-CM pair, but identify the HFC elements between the CMTS and CM to then recover the timing values for those HFC elements.

## **8. What's next?**

In this section, we highlight the next steps for the industry to drive the adoption and use of DTP.

Key to the adoption of DTP is the calibration of the CMTS/RPD/RMD-CM pairs and the HFC elements. With these calibration values collected, a repository where this data is stored needs to be established. CableLabs® is uniquely positioned to be the lead in this calibration initiative.

Having the calibration data available, DOCSIS systems need to introduce automated mechanisms to recover the DTP calibration data of the CMTS-CM pairs, identify the HFC elements between the CMTS-CM pair, and recover the DTP calibration data of those HFC elements.

In parallel, test equipment vendors need to build calibration tools to allow the calibration of CMTS, CM, and HFC elements separately from each other to simplify adoption of DTP.

## 9. Conclusions

The I01 version of the SYNC spec, released on April 20, 2020 describes how to build and deploy an IEEE 1588/PTP participant DOCSIS network. The SYNC specification is targeted at the new and evolving mobile backhaul market, if the DOCSIS system has the required throughput, latency, and timing.

The authors in [2] described how to achieve low latency backhaul over DOCSIS. In terms of timing, the DOCSIS system is already based upon highly precise timing. DTP leverages this asset to provide accurate timing to end applications connected through a CM. Rather than run NTP or PTP over-the-top, the DOCSIS system can be used as-is to generate or correct these timing protocols with a very high degree of precision.

## Abbreviations

|               |                                                 |
|---------------|-------------------------------------------------|
| <b>CM</b>     | Cable modem                                     |
| <b>CMCI</b>   | CM to CPE interface                             |
| <b>CMTS</b>   | Cable modem termination system                  |
| <b>CoMP</b>   | Coordinated Multi Point operation               |
| <b>DOCSIS</b> | Data-Over-Cable Service Interface Specification |
| <b>DTP</b>    | DOCSIS Time Protocol                            |
| <b>DS</b>     | downstream                                      |
| <b>eICIC</b>  | enhanced inter-cell interference coordination   |
| <b>eNB</b>    | Evolved Node B                                  |
| <b>FDD</b>    | Frequency-Division Duplex                       |
| <b>GM</b>     | Grandmaster                                     |
| <b>GNSS</b>   | Global Navigation Satellite System              |
| <b>HetNet</b> | heterogeneous network                           |
| <b>HFC</b>    | hybrid-fiber/coax                               |
| <b>LTE</b>    | Long Term Evolution                             |
| <b>MNO</b>    | mobile network operator                         |
| <b>MSO</b>    | Multiple Systems Operator                       |
| <b>MULPI</b>  | MAC and Upper Layer Protocols Interface         |
| <b>NTP</b>    | network time protocol                           |
| <b>OFDM</b>   | Orthogonal Frequency Division Multiplexing      |
| <b>OTT</b>    | over-the-top                                    |

|             |                         |
|-------------|-------------------------|
| <b>PHY</b>  | physical                |
| <b>PLL</b>  | Phase-Locked Loop       |
| <b>PTP</b>  | Precision Time Protocol |
| <b>RAN</b>  | radio access network    |
| <b>RMD</b>  | Remote MACPHY Device    |
| <b>RPD</b>  | Remote PHY Device       |
| <b>TDD</b>  | Remote PHY Device       |
| <b>TRO</b>  | true ranging offset     |
| <b>TWTT</b> | Two-Way Time Transfer   |
| <b>UE</b>   | user equipment          |
| <b>US</b>   | upstream                |

# Bibliography & References

- [1] J. T. Chapman, "The DOCSIS Timing Protocol (DTP)," in *SCTE Spring Technical Forum*, 2011.
- [2] J. T. Chapman and J. Andreoli-Fang, "Mobile Backhaul over DOCSIS," in *SCTE*, Denver, 2017.
- [3] Cable Television Laboratories, Inc., "Data-Over-Cable Service Interface Specifications; Synchronization Techniques for DOCSIS® Technology," 2020.
- [4] Cable Television Laboratories, Inc., "Data-Over-Cable Service Interface Specifications; MAC and Upper Layer Protocols Interface," 2020.
- [5] J. Andreoli-Fang, J. T. Chapman, T. Liu and D. Poltz, "Blueprint for Mobile Xhaul over DOCSIS," in *SCTE Fall Technical Forum*, 2019.
- [6] J. Andreoli-Fang and J. T. Chapman, "Synchronization for Mobile Backhaul over DOCSIS," in *SCTE Fall Technical Forum*, 2017.

# **Leakage In A High Split World**

## **Detecting and Measuring Upstream Leakage Levels in a One Gpbs Symmetrical High Split Hybrid Fiber Coax Network**

A Technical Paper prepared for SCTE/ISBE by

**John Chrostowski**  
Executive Director  
Comcast  
John\_Chrostowski@comcast.com

**Greg Tresness**  
President  
Arcom Digital, LLC.  
tresness.greg@arcomlabs.com

**Dan Rice**  
VP HFC Architecture Engineering  
Comcast  
Dan\_Rice4@Comcast.com

**Benny Lewandowski**  
Engineering Architect  
Comcast  
Benny\_Lewandowski@Comcast.com

# Table of Contents

| Title                                                                                  | Page Number |
|----------------------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                                   | 4           |
| 2. Capabilities of DOCSIS 3.1 .....                                                    | 4           |
| 3. The Need for High-Split Hybrid Fiber Coax Networks .....                            | 5           |
| 3.1. Impact on Existing Legacy Customer Premise Equipment.....                         | 5           |
| 4. Federal Communications Commission Requirements.....                                 | 6           |
| 5. How Leakage Detection is Currently Measured in the Downstream .....                 | 7           |
| 5.1. Detection of low-level inserted carriers.....                                     | 8           |
| 5.2. Direct QAM detection using correlation processing .....                           | 8           |
| 5.3. Detection of harmonics of OFDM continuous pilots .....                            | 8           |
| 6. Similarities and Differences between Upstream and Downstream Leakage Detection..... | 9           |
| 7. Upstream Leakage Detection Possibilities.....                                       | 10          |
| 7.1. Downstream Out of Band (OOB) CW pilots.....                                       | 10          |
| 7.2. Continuous Waveform Upstream Test Signal.....                                     | 12          |
| 7.2.1. CW Time-Division Multiple Access (TDMA).....                                    | 12          |
| 7.2.2. CW Frequency-Division Multiple Access (FDMA) .....                              | 14          |
| 7.3. OUDP Burst Test Signal (BTS).....                                                 | 17          |
| 8. Receiver Sensitivity analysis of the proposed detection approaches .....            | 19          |
| 8.1. Downstream Out of Band (OOB) CW pilots.....                                       | 19          |
| 8.2. CW Time-Division Multiple Access (TDMA).....                                      | 19          |
| 8.3. CW Frequency-Division Multiple Access (FDMA).....                                 | 21          |
| 8.4. OUDP Burst Test Signal .....                                                      | 22          |
| 9. Using Full Band Capture to Evaluate Potential Aeronautical Band Leakage Issues..... | 23          |
| 10. Probability of Intercept: Moving Vehicle Analysis .....                            | 24          |
| 11. Digital Leakage Detection Test Results .....                                       | 26          |
| 11.1. CW Time-Division Multiple Access (TDMA).....                                     | 26          |
| 11.2. CW Frequency-Division Multiple Access (FDMA).....                                | 29          |
| 11.3. OUDP Burst Test Signal .....                                                     | 30          |
| 12. Conclusions .....                                                                  | 35          |
| Abbreviations.....                                                                     | 36          |
| Bibliography & References .....                                                        | 38          |

## List of Figures

| Title                                                                                                                                                     | Page Number |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| Figure 1 - FCC Requirements Flow Chart .....                                                                                                              | 7           |
| Figure 2 - Leakage test signal injected at 138 MHz, on the channel boundary between adjacent SC-QAM signals (STD channel plan assumed). .....             | 8           |
| Figure 3 - Leakage test signal injected at 120 MHz and 126 MHz, on the channel boundaries below and above SC-QAM signal (STD channel plan assumed). ..... | 8           |
| Figure 4 - Harmonics of OFDM Continuous Pilots .....                                                                                                      | 9           |
| Figure 5 - Example HFC Network Downstream to Upstream Level Deltas .....                                                                                  | 10          |
| Figure 6 - High Split RPD/Node with OOB Filtering .....                                                                                                   | 11          |
| Figure 7 - High Split Upstream Channel Configuration with OOB.....                                                                                        | 11          |
| Figure 8 - Out of Band Band Pass Filter Response with OOB Signal and Leakage Tones .....                                                                  | 11          |
| Figure 9 - Illustration of permanent generation of CW test bursts by CM.....                                                                              | 13          |

|                                                                                                                                                             |    |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| Figure 10 - Illustration of CW test signals within exclusion BW of OFDMA signal (4k FFT mode).....                                                          | 13 |
| Figure 11 - Illustration concept of continuous generation CW tones by all CMs .....                                                                         | 15 |
| Figure 12 - Generation of OUDP Test Bursts in a Service Group.....                                                                                          | 18 |
| Figure 13 - Patterns of OUDP Test Burst in Time-Frequency Domain .....                                                                                      | 19 |
| Figure 14 - FBC Capture Showing No Ingress in FM Band.....                                                                                                  | 23 |
| Figure 15 - FBC Showing Ingress in the FM Band.....                                                                                                         | 24 |
| Figure 16 - Probability of Intercept (POI), Detection of DS and US leaks from moving vehicle .....                                                          | 24 |
| Figure 17 - Distance covered in 1 second versus vehicle speed and ratio of the minimum distance to the maximum distance a vehicle could have travelled..... | 25 |
| Figure 18 - Leakage Detection OUDP Test Bursts and Data Transmission for a SG.....                                                                          | 26 |
| Figure 19 - CW Leakage Detection Test Setup .....                                                                                                           | 27 |
| Figure 20 - CW test burst in time domain and spectrum of CW bursts .....                                                                                    | 27 |
| Figure 21 - Spectrum of CW test bursts within exclusion BW of OFDMA signal .....                                                                            | 28 |
| Figure 22 - FFT spectrums and results of leak detection for boosting gain K=+10 dBc.....                                                                    | 28 |
| Figure 23 - Continuous CW leakage Detection Test Setup.....                                                                                                 | 29 |
| Figure 24 - The Spectrum of Groups of CW tones within OFDMA exclusion BW .....                                                                              | 29 |
| Figure 25 - Leakage Detector settings for detection of CW tones and minimum detected signal .....                                                           | 30 |
| Figure 26 - Test bench block diagram .....                                                                                                                  | 31 |
| Figure 27 - Pattern of the OUDP test burst in time – frequency domain .....                                                                                 | 31 |
| Figure 28 - Pilot's BPSK modulation in OUDP test burst in Figure 27 [1] .....                                                                               | 32 |
| Figure 29 - Allocation of OUDP test burst in frequency domain of aeronautical band .....                                                                    | 32 |
| Figure 30 - Arbitrary waveform generation settings at DLCG for simulation OUDP test burst.....                                                              | 33 |
| Figure 31 - Arbitrary waveform transmission settings at DCLG for simulation OUDP test burst.....                                                            | 33 |
| Figure 32 - OUDP test burst in time and frequency .....                                                                                                     | 34 |
| Figure 33 - Response of matched filter for OUDP test burst with pilot pattern 11. ....                                                                      | 34 |
| Figure 34 - Results of implementation of UUDP burst detector into field meter .....                                                                         | 35 |

## List of Tables

| <b>Title</b>                                                         | <b>Page Number</b> |
|----------------------------------------------------------------------|--------------------|
| Table 1 - CFR § 76.605 Signal Leakage Limits .....                   | 6                  |
| Table 2 - Relative level of CW tones vs OFDMA subcarriers, dBc ..... | 17                 |
| Table 3 - Estimated Sensitivity for Different Boosting Gains, K..... | 20                 |
| Table 4 - Sensivity for Pilot Pattern 11 .....                       | 23                 |
| Table 5 - Summary of Leakage Detection Methods .....                 | 36                 |



## 1. Introduction

System leakage monitoring is an integral and extremely important aspect of system maintenance. Federal Communications Commission (FCC) leakage requirements, test and mitigation methods for present cable systems with 5-42 MHz and 5-85 MHz return bands are well understood. Leakage signal sources, monitoring equipment, methods and measurement programs are all in place for measurements in the downstream band, 54-1000 MHz, with an emphasis on the aeronautical band of 108-137 MHz. The FCC sets maximum individual signal leakage levels for cable systems, with more stringent limits on cable systems that may interfere with aeronautical and navigation communications.

In a traditional cable system with a 5-42 MHz return band, high level signals in the downstream are present at the headend or node output, and at the home, the downstream signal is at its lowest level. What happens if you increase the upper boundary of the upstream signal path to 204 MHz or higher? The system is essentially turned upside down, with the highest signal levels in the aeronautical band (108-137 MHz) at the home, and the lowest upstream signal levels at the headend or node.

This paper will look at the implications of these inverted plant levels. How do the FCC's cable signal leakage requirements apply to this scenario? Can the same leakage tools and methods be used? What changes and considerations need to be made? These questions not only apply to high split plants with 204 MHz upstream, but also to future full duplex data over cable service interface specification.

(FDX DOCSIS) and extended spectrum DOCSIS (ESD) systems, each with upstream signals well beyond 204 MHz.

This paper will review these considerations and discuss the implications of the “inverted” plant. Different options and scenarios will be examined, and their implementation and feasibility evaluated to help readers ensure a solid leakage measurement program is in place for when these plant upgrades are implemented.

## 2. Capabilities of DOCSIS 3.1

DOCSIS® 3.1 (D3.1) specifications expand the useable frequency spectrum as well as provide a more efficient use of that spectrum compared to its predecessor DOCSIS®3.0 (D3.0.) With downstream operation out to 1218 MHz along with a high-split (204/258 MHz) option, network capacity can be dramatically improved, particularly in the upstream direction. Included in these exceptional capabilities are orthogonal frequency-division multiplexing (OFDM) channels in the downstream and orthogonal frequency-division multiple access (OFDMA) channels in the upstream. OFDM and OFDMA channels permit a wide-ranging encompassed spectrum-improving capacity, while still allowing very small minislots for enhanced bandwidth utilization. With the 204 MHz upstream spectrum available, and with optimal use of the OFDMA technology, 1 Gbps symmetrical services become possible. In addition, higher-order modulations, including 1024 to 4096 QAM downstream and 256 to 2048 QAM upstream, as well as enhanced forward error correction (FEC) methods in both directions, provide another example of how D3.1 can improve overall capacity and customer experience. Modifications to the hybrid fiber coax (HFC) network parameters, as they pertain to a D3.1 upgrade, need to be continuously evaluated for impact on legacy and future FCC compliance.

### 3. The Need for High-Split Hybrid Fiber Coax Networks

The demands of an HFC network, and, in particular network traffic, continue to increase. Downstream traffic has seen tremendous growth, with over-the-top (OTT) video driving most of the surge. With D3.1, the downstream network is more than capable of handling this traffic increase. A typical High-Split D3.1 downstream channel lineup can accommodate over ninety (256 QAM) D3.0 channels and two 192 MHz (4096 QAM) OFDM D3.1 channels, exceeding throughputs of 9 Gbps.

The limited 5-42 MHz signaling spectrum is typically where the bottleneck in the upstream network performance occurs. Applications and hardware are continually being developed that utilize the upstream more frequently, and with higher bandwidth needs. Internet of Things (IOT), video doorbells, security cameras, online gaming and cellphone Wi-Fi handoffs are all taxing the upstream network.

Currently, and even more relevant, is the impact that COVID-19 is having on the network. In three months, from early March to late May, multiple system operators (MSOs) saw a 32% increase in upstream traffic and an 11% increase in downstream traffic. This added growth is due to a sharp rise in video conferencing and video streaming, now that so many employees and students are working/learning from home.

Beyond D3.1 improvements in efficiency, node splits have been the go-to method to increase upstream capacity. While node splits do increase capacity, they do so without increasing the upstream peak data rate. Upgrading to a mid-split (5-85 MHz) will increase the upstream data throughput to ~ 500 Mbps, and upgrading to high-split (5-204 MHz) will allow >1 Gbps upstream throughput, future proofing the network and keeping the DOCSIS network competitive with alternative technologies. Full duplex (FDX) technology further expands the upstream capabilities, enabling multi-Gbps upstream speeds.

High-split D3.1 technology and products are in high demand and readily available. End-to-end solutions from cable modem termination system (CMTS) all of the way through the plant to the customer premise equipment (CPE) can be procured and installed. This allows operators to upgrade to a high-split network quickly, drastically improving upload speeds, while maintaining extraordinary downstream traffic capacity.

#### 3.1. Impact on Existing Legacy Customer Premise Equipment

A large majority of cable modems deployed in the field do not support high-split operation and will operate “business as usual” with the channels that reside in their existing spectrum allocations. Moreover, legacy set-top boxes (STBs) are limited to control channel frequencies between 70 and 130 MHz in the downstream direction, so the out-of-band (OOB) channel will not pass through any duplex filters notched at 204/258 MHz that are used in the traditional HFC network. MSOs are faced with two options: Replace all legacy STBs to support high-split systems with DOCSIS set-top gateways (DSGs) (with the OOB control channel carried over DOCSIS) or enable the OOB control channel signal to propagate through the high-split network to the legacy STB in the upstream spectrum band. For the latter solution, signal requirements defined by SCTE 55-1, 55-2 and SCTE 40 must be maintained, and with a broad range of possible OOB frequencies, challenges from a scalability and standardized solution perspective will need to be resolved.

## 4. Federal Communications Commission Requirements

Understanding the FCC requirements for leakage is a daunting task. The FCC requirements are a maze of different related and unrelated requirements. We'll take you through the main requirements as they relate to the CATV network today and as they relate to the high split network.

There are a few main key FCC requirements. **47 CFR § 76.605** Technical Standards defines the CATV signal and signal leakage limits for both analog and digital signals. Leakage limits are defined in § 76.605 and are shown below in Table 1.

**Table 1 - CFR § 76.605 Signal Leakage Limits**

| Frequencies                                                      | Signal leakage limit | Distance in meters (m) |
|------------------------------------------------------------------|----------------------|------------------------|
| Analog signals less than and including 54 MHz, and over 216 MHz  | 15 $\mu$ V/m         | 30                     |
| Digital signals less than and including 54 MHz, and over 216 MHz | 13.1 $\mu$ V/m       | 30                     |
| Analog signals over 54 MHz up to and including 216 MHz           | 20 $\mu$ V/m         | 3                      |
| Digital signals over 54 MHz up to and including 216 MHz          | 17.4 $\mu$ V/m       | 3                      |

The following FCC requirements are applicable to all multichannel video programming distributors (MVPDs): §76.605(d) Technical Standards, § 76.611 Cable television basic signal leakage performance criteria, § 76.612 Cable television frequency separation standards, § 76.613 Interference from a multichannel video programming distributor (MVPD), 76.614 Cable television system regular monitoring, 76.616 Operation near certain Aeronautical and Marine Emergency Radio Frequencies, 76.617 Responsibility for interference, 76.1803 Signal Leakage monitoring and 76.1804 Aeronautical frequencies: leakage monitoring (CLI)

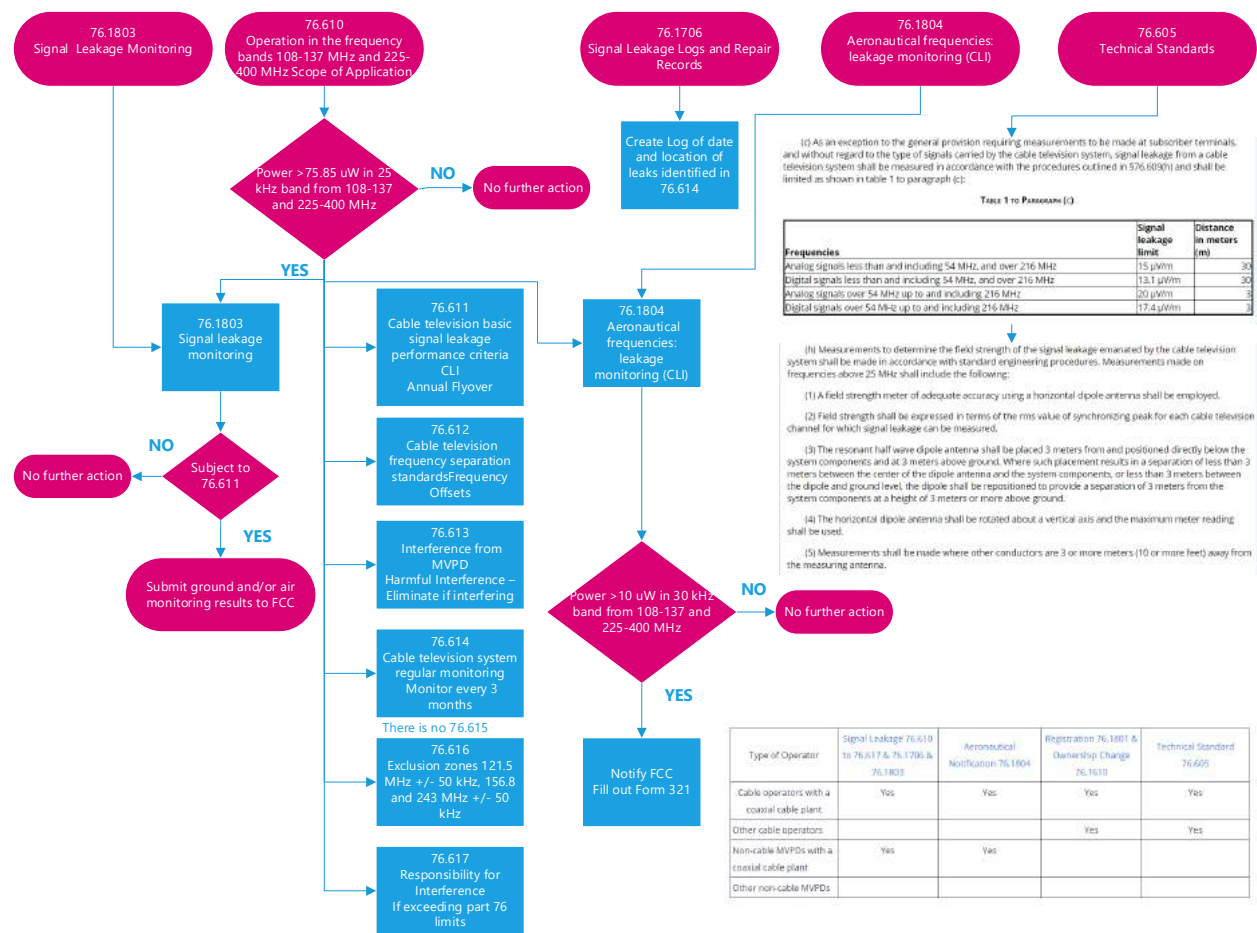
With respect to 47 CFR § 76.610 – Operation in the frequency bands 108-137 MHz and 225-400 MHz - scope of application, this requirement states: The provisions of §§ 76.605(d), 76.611, 76.612, 76.613, 76.614, 76.616, 76.617, 76.1803 and 76.1804 are applicable to all MVPDs (cable and non-cable) transmitting analog carriers or other signal components carried at an average power level equal to or greater than 100 microwatts across a 25 kHz bandwidth in any 160 microsecond period or transmitting digital carriers or other signal components at an average power level of 75.85 microwatts across a 25 kHz bandwidth in any 160 microsecond period at any point in the cable distribution system in the frequency bands 108-137 and 225-400 MHz for any purpose.

The requirements above warrant a look at the upstream levels generated at the cable modem, to determine if the additional sections referenced apply in a High Split system. Take, for example, an upstream channel configuration with four SC-QAM channels in the 5-42 MHz band and two OFDMA channels, each with a bandwidth of 81 MHz, from 42-204 MHz.

Per the D3.1 specifications, the maximum transmit power from a cable modem is 65 dBmV. With 4 SC-QAMs each 6.4 MHz wide and two OFDMA channels from 42-204 MHz, the total upstream utilized bandwidth is 187.6 MHz. This equates to a power per 25 KHz bandwidth of 5.6 uW and a power per 30 KHz bandwidth of 6.7 uW. This transmit power is below the 75.85 uW threshold in 76.610, which triggers applicability to; §§ 76.605(d), 76.611, 76.612, 76.613, 76.614, 76.616, 76.617, 76.1803 and 76.1804. Even though 76.1804 is not applicable, 76.1804 had a different threshold of 10 uW. In this example, the threshold of 10uW in 76.1804 is also not exceeded. 76.616 is also not applicable per the requirements, but the exclusion zones required to avoid any impact to these emergency beacons are supported by D3.1, are minimal in bandwidth impact, and have minimal effect on upstream capacity or peak speeds.

D3.1 has a maximum power limit for OFDMA channels of  $\leq 24$  MHz of  $-9$  dBmV/Hz. As another example, consider a single OFDMA channel of 24 MHz, which happens to overlap with the 108-137 MHz aeronautical band. While a deployment that only uses 24 MHz is not a productive use of spectrum and is unlikely to be deployed, it could possibly be done as the worst case power spectral density, per the D3.1 specifications. The power per Hz is constant at  $-9$  dBmV. This equates to a power per 25 KHz bandwidth of 42 uW, and a power per 30 KHz bandwidth of 50.4 uW, which is again below the threshold of 75.85 uW in 76.610, which triggers applicability to §§ 76.605(d), 76.611, 76.612, 76.613, 76.614, 76.616, 76.617, 76.1803 and 76.1804. 76.1804 has a power level requirement of 10 uW in a 30 KHz band. The threshold of 76.1804 is exceeded and FCC notification on form 321 is required. The same comments on 76.716 apply as in the first example.

To help navigate the FCC requirements, Figure 1 shows the majority of requirements and applicable conditions which apply to CATV systems.




**Figure 1 - FCC Requirements Flow Chart**

## 5. How Leakage Detection is Currently Measured in the Downstream


There are several methods of signal leakage detection utilized by operators in today's HFC network. These can be broadly broken down into three categories: detection of low-level inserted carriers, direct QAM detection using correlation processing, and detection of harmonics of OFDM continuous pilots.

### 5.1. Detection of low-level inserted carriers

With this technique, two CW carriers are inserted into the network, at either the hub or at the R-PHY or MAC-PHY node, and the leakage signal is captured using a fast fourier transform (FFT) detector. The carriers are typically configured to be -30dBc, relative to the SC-QAM digital channel power. Various leakage detection vendors utilize different configurations of CW carrier placement or configuration of the inserted signal. Some examples of widely used signal types are provided in Figure 2 and Figure 3.



**Figure 2 - Leakage test signal injected at 138 MHz, on the channel boundary between adjacent SC-QAM signals (STD channel plan assumed).**




**Figure 3 - Leakage test signal injected at 120 MHz and 126 MHz, on the channel boundaries below and above SC-QAM signal (STD channel plan assumed).**

### 5.2. Direct QAM detection using correlation processing

With this technique, downstream signal samples are captured at the headend at the desired leakage detection frequency. The signal samples are time stamped using a GPS reference clock and transmitted to a leakage detector in the field. The leakage detector receives signals leaking from the HFC network, time stamps the received signals, and then applies a cross-correlation process on the two signal sets to resolve the detected leak.

### 5.3. Detection of harmonics of OFDM continuous pilots

With this technique, an FFT detector is again utilized in the field meter to capture the leakage signal – but with this method no additional tagging or CW signal is inserted into the HFC network. The detected signals are the existing harmonics of the OFDM continuous pilots, as show below in Figure 4.



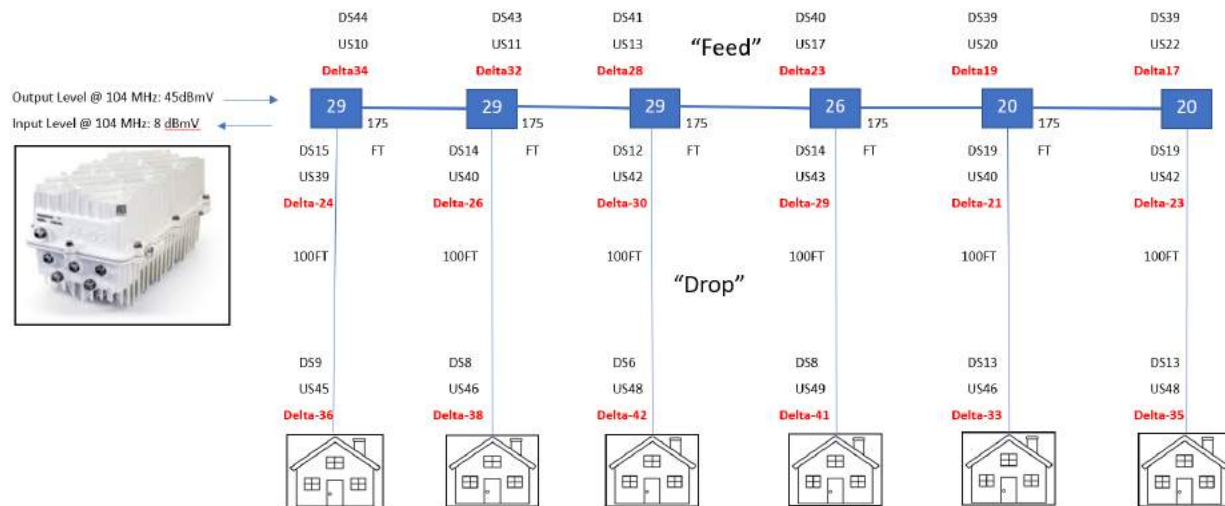
**Figure 4 - Harmonics of OFDM Continuous Pilots**

Unfortunately, none of the status quo approaches are suitable for use in the high-split network. This is mainly because the signals coming from the cable modem (CM) are inherently burst-like in nature, and instead of the leakage detection signal being generated from one source, as is the case with downstream detection, the detection signal is generated from the numerous non-coherent CMs in the upstream network. There are, however, some alternatives, presented below.

## 6. Similarities and Differences between Upstream and Downstream Leakage Detection

For upstream signal generation, the high-split cable modem must be the source of the leakage detection signal. With the majority of the emphasis for downstream leakage detection on the “feed” side of an HFC network, attention now inverts to the “drop” side, including the home -- where upstream signal levels in a high-split network now encompass the aeronautical band and are at much higher levels than the downstream. Cable modems can transmit a maximum power of 65 dBmV for a single 6.4 MHz SC-QAM channel and receive downstream signals within a recommended range of (10 to -10 dBmV, 6 MHz BW.) In this instance, the leakage level delta in terms of uV/m can be over **2000** times higher in the upstream direction versus the downstream, if the leak is on the output of the cable modem. Having multiple cable modems transmitting leakage tones (per the technique described in Section 5.1) at the same time and frequency can cause constructive and destructive signal combining because of the different phases of the signal. This will create a high likelihood of inaccurate leakage detection measurements. Furthermore,

upstream leakage detection signals should be in burst mode operation, as opposed to continuous mode downstream tones, to limit the impact on upstream capacity and to protect the upstream power budgets of the optical components. Figure 5 illustrates the inverse relationship of DS to US signal levels between the “feed” and “drop” side of a tap.



**Figure 5 - Example HFC Network Downstream to Upstream Level Deltas**

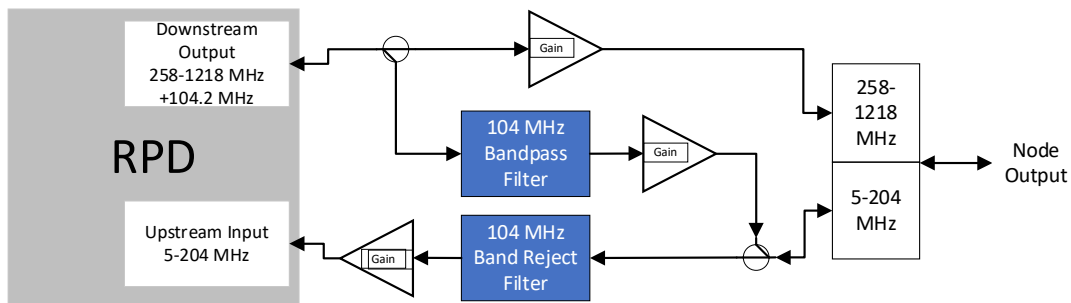
## 7. Upstream Leakage Detection Possibilities

### 7.1. Downstream Out of Band (OOB) CW pilots

With millions of legacy set top boxes that utilize SCTE 55-1 and SCTE 55-2 out-of-band data carriers in the forward/downstream, this functionality can be maintained through the distributed access architecture (DAA) node and amplifiers in high split systems. Two of the typical frequencies for the OOB carriers are 75.25 MHz and 104.2 MHz in 5-42 MHz systems. In mid split systems with 5-85 MHz upstream, 104.2 MHz is the typical frequency for the OOB carrier in the downstream.


In a high split system, 104.2 MHz sits in the middle of the upstream band and is also very close to the aeronautical band. To utilize a 104.2 MHz OOB carrier in a high split system, the 104.2 MHz OOB signal is generated in the CMTS or remote physical device (RPD), and via couplers and filtering this signal can be injected onto the plant on the low side of the diplex filter. Notch filters are required in the upstream to prevent the OOB signal from entering the return path receiver of the CMTS or RPD at a level that would cause distortion or interference. See Figure 6 for a simplified OOB implementation example in a high split RPD/node.

## High Split Node with Out of Band




**Figure 6 - High Split RPD/Node with OOB Filtering**

Leakage tones can be generated in the RPD and placed close to the OOB carrier, and still fit within the 104 MHz bandpass filter passband. In a 5-204 MHz upstream, the two OFDMA channels can be placed above and below 104.2 MHz, to create a band gap for the 104.2 MHz OOB carrier. See Figure 7.



**Figure 7 - High Split Upstream Channel Configuration with OOB**

The 104 MHz bandpass filter can be made wide enough to include both the 104.2 MHz OOB and leakage tones. The leakage tones will then proceed down the plant along with the OOB carrier. See Figure 8.



**Figure 8 - Out of Band Band Pass Filter Response with OOB Signal and Leakage Tones**

Although the leakage tones can be generated close to the lower edge of the aeronautical band, all the difficulties of the inverted plant and level differences between the downstream and upstream levels will make it difficult to correlate leakage measurements. Looking at Figure 5, the downstream-to-



upstream signal level delta is a positive 37 dB, and slowly decreases as you travel down the hard line plant, to a delta of 17 dB at the input to the last tap. Once you move to the drop side of the tap, the downstream-to-upstream delta decreases dramatically, with the largest negative delta at the premises. A method could be developed, based on plant maps and GPS coordinates, as leakage is measured along the plant. The maximum CM transmit level of 65 dBmV could be assumed, or additional information on CM transmit levels could also be determined from monitoring tools.

High Split devices will be introduced into the plant incrementally, as needed. Using the leakage tones generated at the RPD or CMTS could generate false readings, as the leakage tones are distributed to the whole plant and all homes. In most instances, using the default transmit value of 65 dBmV will be higher than the actual CM transmit values, and potentially create false high leakage readings. With the extensive mapping and software updates required, and the difficulty in correlating downstream-to-upstream levels, the other approaches described in this paper are likely more practical to implement. The goal is to be able to measure leakage events as reliably as possible, to meet regulatory requirements, without unnecessary increases in operational expense, or unnecessary risks for compliance issues.

## **7.2. Continuous Waveform Upstream Test Signal**

One of the possibilities for upstream leakage detection is for the CM to generate continuous wave (CW) carriers. The CW carriers generated by the CMs cannot be continuously and simultaneously present at any one frequency, because the numerous out-of-phase signals would combine randomly and, as a direct result, yield random detection results. As an alternative, two CW approaches are presented below -- one with the CW burst controlled in time by the CMTS, and the second where each CM in a node transmits at a unique frequency. One important obstacle that will need to be overcome in these two approaches involves the fact that both are outside of the existing D3.1 specification. Also, current CMs are not able to generate the required CW carriers without updated firmware, and face constraints without updated hardware.

### **7.2.1. CW Time-Division Multiple Access (TDMA)**

In this scenario, there is continuous cyclical generation of two CW test bursts by CMs, under CMTS control, in an exclusion segment of bandwidth of the OFDMA spectrum within the aeronautical band. Detection of the leakage signal is provided by advanced FFT overlap processing, with leak signal validation accomplished by measuring the frequency offset between CW test tones, and by verifying the time sync and duration of the CW bursts. Figure 9 and Figure 10 below describe the burst schema and carrier placement.

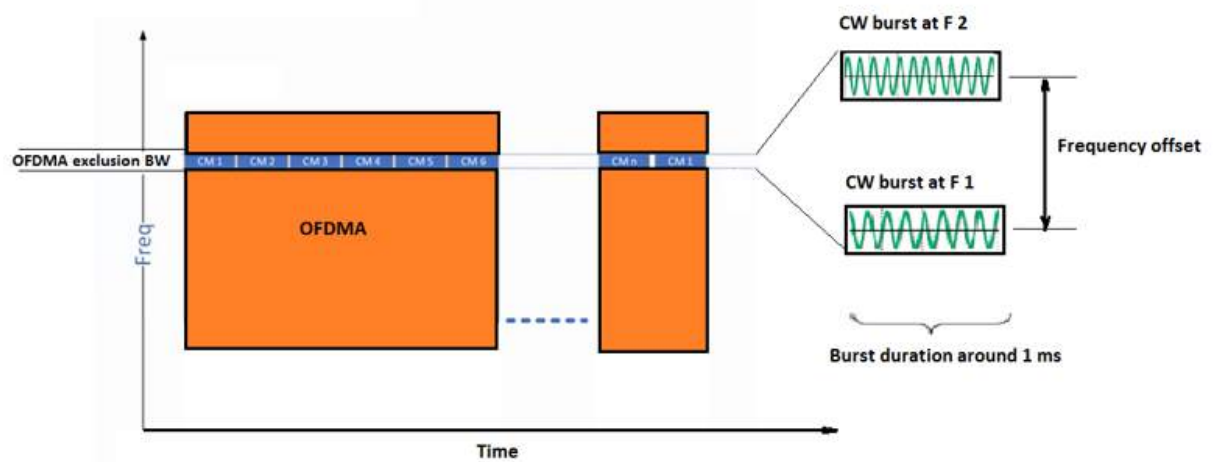


Figure 9 - Illustration of permanent generation of CW test bursts by CM




Figure 10 - Illustration of CW test signals within exclusion BW of OFDMA signal (4k FFT mode)

#### 7.2.1.1. Analysis of the parameters of the CW test signal:

The following parameters of the CW test signal are relevant for leakage detection and for compatibility with HFC data transmission:

- Frequency location of the CW test signal;
- Duration of the CW burst;
- Frequency offset between the CW bursts central frequencies;
- Level of the CW test signal relative to the level of the OFDMA subcarriers;
- **The frequency location of the CW test signal** should be within or very close to aeronautical band 108 - 137 MHz for accurate CLI reporting. Given that 108 MHz is the

edge frequency for FDX, it makes sense to assume that high split HFC upstream OFDMA spectrum also will also occupy the 108 MHz to 204 MHz band. This means the CW test signal could be placed near or within the OFDMA spectrum, in an exclusion bandwidth (BW), or below the edge frequency of 108 MHz, at a guard bandwidth within the lower adjacent OFDMA upstream signal.

- **The duration of the CW test burst** should be as long as possible, to provide improved sensitivity of leakage detection. On the other hand, the duration of the CW test burst is limited by the condition that each CM in a node serving group should be granted a transmit (Tx) test burst at least two times during the 1 second leak measuring session. Therefore, for a typical number of CMs in node, from 100 to 500 CMs, the maximum duration of the CW test burst is limited to around 1 ms.
- **The frequency offset between CW tones** should be as low as possible, to minimize the exclusion BW area of the OFDMA spectrum. On the other hand, the frequency offset needs to be sufficient to provide the FFT analyzer with enough resolution for correct leak signal validation. The bandwidth of a CW burst with a duration of 1 ms is a few kHz. Therefore, the minimum frequency offset should be around 10 kHz or more. To provide the minimum interfering effect, when receiving the OFDMA signal at the CMTS, it makes sense to establish the central frequencies of the CW bursts equally to the central frequencies of OFDMA subcarriers. In this case, the frequency offset between CW bursts will be 25 kHz (4k FFT) or 50 kHz (for 2k FFT) mode. The minimal exclusion bandwidth for placing CW test signals will be two subcarriers, however additional guard band carriers may be required for the exclusion area.
- **The level of each CW test signal** should be boosted as much as possible, relative to the OFDMA subcarriers level, to provide better sensitivity of leak detection. On the other hand, the level of the CW test burst must be below the possible interfering threshold at the receive (Rx) side of CMTS. According to the D3.1 specifications, OFDMA pilots could be boosted by at least 4.7 dB:[1]

The CM MUST boost pilots and complementary pilots by a factor of 3 in power (about 4.7 dB).

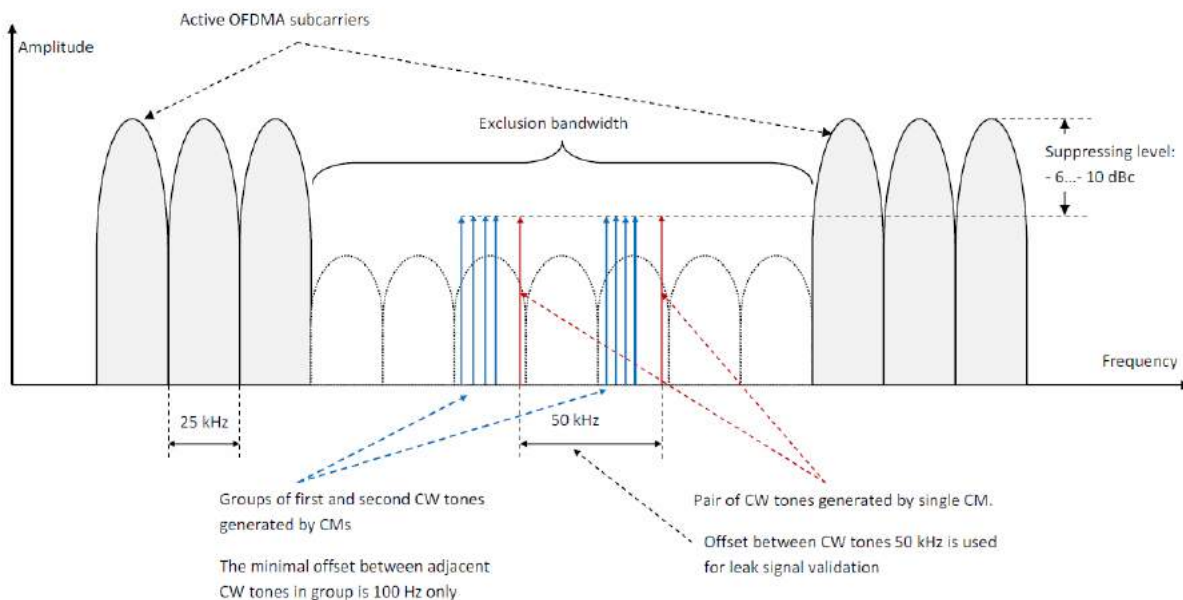
Another approach to boosting the gain could be to increase the exclusion BW in the OFDMA spectrum and add the energy of the excluded subcarriers to the CW test signal. For example, in case of exclusion BW = 1.6 MHz, the boosting gain could be increased to +15 dBc. In this scenario, by placing the CW test signal outside the OFDMA spectrum (for example, below 108 MHz), the boosting gain could be increased up to + 20 dBc, where the power of CW test signal will be equal to the energy of one upstream SC-QAM. As a rule of thumb, the larger the spectrum exclusion, the more lost capacity for user data. Capacity calculation and OFDMA testing indicates that the goal of a 1 Gbps US product capability will require precise attention to spectral efficiency.

### 7.2.2. CW Frequency-Division Multiple Access (FDMA)

This approach requires continuous generation of a pair of CW tones by each CM. The frequency offset between CW tones is the same at all CMs used for leak signal validation. The central

frequencies of the CW tone pairs at different CMs within one node will have a specific frequency offset, with each CM in a node transmitting at a unique frequency. This frequency offset is required for frequency resolution of signals from different CMs in the leak detector, and for correct measurement of the leak level, such as in the scenario where there is combining of CW tones from many CMs in the trunk line. All CW tones are located within the exclusion bandwidth of the OFDMA spectrum within the aeronautical band. The level of CW tones is suppressed relative to the level of the OFDMA subcarriers, to prevent interfering with the CMTS Rx side. Detection of the leak signal is provided by FFT processing. This is the same technology as is currently used for detection of DS harmonics of OFDM continuous pilots, and for pilots inserted between SC-QAM channels.

With this approach, permanent control from the CMTS side is not required, as is the case of CW time division multiple access (TDMA) burst or OFDMA upstream data profile (OUDP) approaches. The generation of CW tones can be activated at the CMs during initial provisioning by using direct access via telnet, simple network management protocol (SNMP) or the web. This provisioning of tone frequency, and the frequency accuracy, calibration to OFDMA power levels, and fidelity is not supported in current standards and would need to be agreed upon by all high split CM vendors, through the CableLabs specification and certification process. Figure 11 describes this concept.



**Figure 11 - Illustration concept of continuous generation CW tones by all CMs**

### **7.2.2.1. Analysis of parameters of CW tone signal:**

The following parameters of the CW test signal are relevant for leakage detection and for compatibility with HFC data transmission:

- Frequency location of CW signal;
- Frequency offset between CW tones from different CMs within one node;
- Frequency offset between CW tones in each CM for leak validation;
- Level of CW tones relative OFDMA subcarriers and reasonable exclusion bandwidth.

- **The frequency location of CW tones** should be within or very close to the aeronautical band 108 - 137 MHz for accurate FCC reporting. Given that 108 MHz is the edge frequency for FDX, it makes sense to assume that for high split HFC, the upstream OFDMA spectrum also will be from 108 MHz to 204 MHz. This means that the CW tone signal could be placed within the OFDMA spectrum at some exclusion BW, or below the edge frequency of 108 MHz at the guard band of the lower adjacent OFDMA upstream signal.
- **The frequency offset between CW tones from different CMs within node** should be big enough to provide good frequency resolution of tones from different CMs in the leak detector FFT processor for correct measurement of leak levels. Absent this frequency offset, tones from different CMs will add with different phases, and the leak signal within the hardline will have random fluctuating levels, depending upon the number of accumulated CW tones. On the other hand, the frequency offset should be as minimal as possible, to reduce the required exclusion bandwidth in OFDMA spectrum. The frequency resolution of current FFT leak detectors is below 10 Hz, so a minimum frequency offset could be selected as 100 Hz. For a typical number of 250 CMs in a node, the CW tones will be grouped in two bandwidths (see Fig.10) of 25 kHz each within one OFDMA subcarrier.
- **The frequency offset between CW tones in each CM is used for leak signal validation.** This offset should be two times more than the bandwidth occupied by each group of CW tones (see Figure 11) to prevent overlapping groups of CW tones in the frequency domain. Therefore, the offset depends both upon the minimum offset and number of CMs in the node. For minimal frequency offset = 100 Hz and the number of CMs in node equal to 250 the minimal offset for leak validation will be 50 kHz as shown in Fig.10.
- **The level of each CW tone** should be as high as possible to provide good sensitivity of leak detection. On the other hand, the energy of CW tones from all CMs in node will be combined at the CMTS Rx side. This means that the level of CW tones should be limited (suppressed), relative to the level of the OFDMA subcarrier, to prevent interfering with and overloading the CMTS. It should be noted that increasing the excluded BW allows an increased level of the CW tones, from a point of view of saving the total energy of the signal in the same BW. Additionally, according to the DOCSIS 3.1 specs, OFDMA pilots can be boosted by at least 4.7 dB:[1]

The CM MUST boost pilots and complementary pilots by a factor of 3 in power (about 4.7 dB).

With this, we can assume the energy of all CW tones must be no more than energy of all excluded subcarriers, plus 4.7 dB. Based on this assumption, the relative level of one CW tone is defined by formula:

$$L \text{ (dBc)} = 10 \text{ Log } (M / 2K) + 4.7 \text{ dB},$$

where:

M is number of excluded subcarriers;

K is maximum number of CMs in node.

Table 2 below shows the relative level of CW tones dependent on the exclusion zone and the number of CMs. Continuously dedicated spectral exclusion bands of multiple MHz are material in the challenges they present to achieve the goal of Gbps US services.


**Table 2 - Relative level of CW tones vs OFDMA subcarriers, dBc**

| Exclusion BW     | Number of CMs in node |       |        |        |        |
|------------------|-----------------------|-------|--------|--------|--------|
|                  | 100                   | 200   | 300    | 400    | 500    |
| <b>400 kHz</b>   | -6.27                 | -9.28 | -11.04 | -12.29 | -13.26 |
| <b>800 kHz</b>   | -3.26                 | -6.27 | -8.03  | -9.27  | -10.25 |
| <b>1,600 kHz</b> | -0.25                 | -3.26 | -5.02  | -6.26  | -7.24  |
| <b>2,000 kHz</b> | +0.72                 | -2.29 | -4.05  | -5.03  | -6.27  |

### 7.3. OUDP Burst Test Signal (BTS)

This leakage detection method uses an OFDMA Upstream Data Profile (OUDP) burst that is generated by each high split cable modem, which is used to detect and monitor leakage in the aeronautical band of a high-split (204/258 MHz) HFC network. Within the DOCSIS specification, there are several predefined OUDP pilot patterns. Existing pilot pattern 11 contains the densest concentration of pilots, and as such it makes sense to use exactly this pattern for the CM burst signal used for detection. Detection of the signal is realized by utilizing a matched filter<sup>1</sup> for the predefined pilot pattern. Scheduling and overall configuration of the cable modem OUDP burst signals will be done through the CMTS service gateway (SGW.) OUDP burst test signal (BTS) has several advantages compared to the previous solutions, the most noteworthy being not having to modify existing D3.1 specifications for cable modem upstream signal generation requirements, and not having to update firmware in existing CMs to add the capability of generating CW carriers. Because there will initially be a limited number of high split cable modems, this approach is also advantageous in that it accommodates the overall time needed to generate the OUDP test bursts. Even in a larger node, sufficient time will exist for data to be interleaved into that spectrum, as needed. In the future, when all modems are high split or FDX capable, the OUDP test bursts can be scheduled only when needed, freeing that spectrum for data bursts most of the time.

An illustration of this OUDP BTS approach, with one frame of 8 symbols and 4 minislots, is shown **Figure 12**.



**Figure 12 - Generation of OUDP Test Bursts in a Service Group**

Using an OFDMA OUDP burst allows configurability in optimizing frequency placement in the upstream band; minimizes any impact to the overall upstream bandwidth/throughput; and optimizes duration to maximize the sensitivity of the receiver. The OFDMA OUDP burst is able to cycle through all the cable modems on the node and have a good probability of intercept (POI) for leakage measurements, as detailed in Section 10 Probability of Intercept: Moving Vehicle Analysis.

One example of an OUDP Burst Signal which can be configured in a high split cable modem is detailed below and shown in Figure 13. The parameters defining the OUDP signal are used to generate the matched filter within the leakage detector:

- Symbols Per Frame (K) = 6 (may be modified to improve network efficiency and robustness)
- Modulation Order = 64 QAM
- Pilot Pattern = 11
- Center Frequency of OUDP Signal = 136.0125 MHz
- 4 Minislots (1.6 MHz Upstream Bandwidth with the 4 adjacent minislots to the center frequency above)
- Number of Frames = 8 and 2.16 mS in transmit time duration. (For 256 House Holds Passed total roundtrip time = 552.96 mS)
- Transmit Power equal to the surrounding OFDMA P= 1.6 MHz channel transmit power
- 4K Fast Fourier Transform (FFT) = 40 uS per symbol + Cyclic Prefix
- Cyclic Prefix = 5.0 uS (may be modified to improve network efficiency and robustness)



- Window Roll off Period = 0.9375  $\mu$ S (may be modified to improve network efficiency and robustness)

The pattern described above and illustrated in the far right of Figure 13 utilizes 8 frames of OUDP pilot pattern 11. This configuration provides the most pilot energy within the OUDP burst signal, which results in optimized sensitivity for detection, as compared to the other variants shown below. If there was the ability to define a new OUDP pilot pattern within the DOCSIS spec that contained an even more dense configuration of pilots, that would be a more spectrally efficient approach, yielding improved sensitivity.

Variants of OUDP test bursts for increasing sensitivity

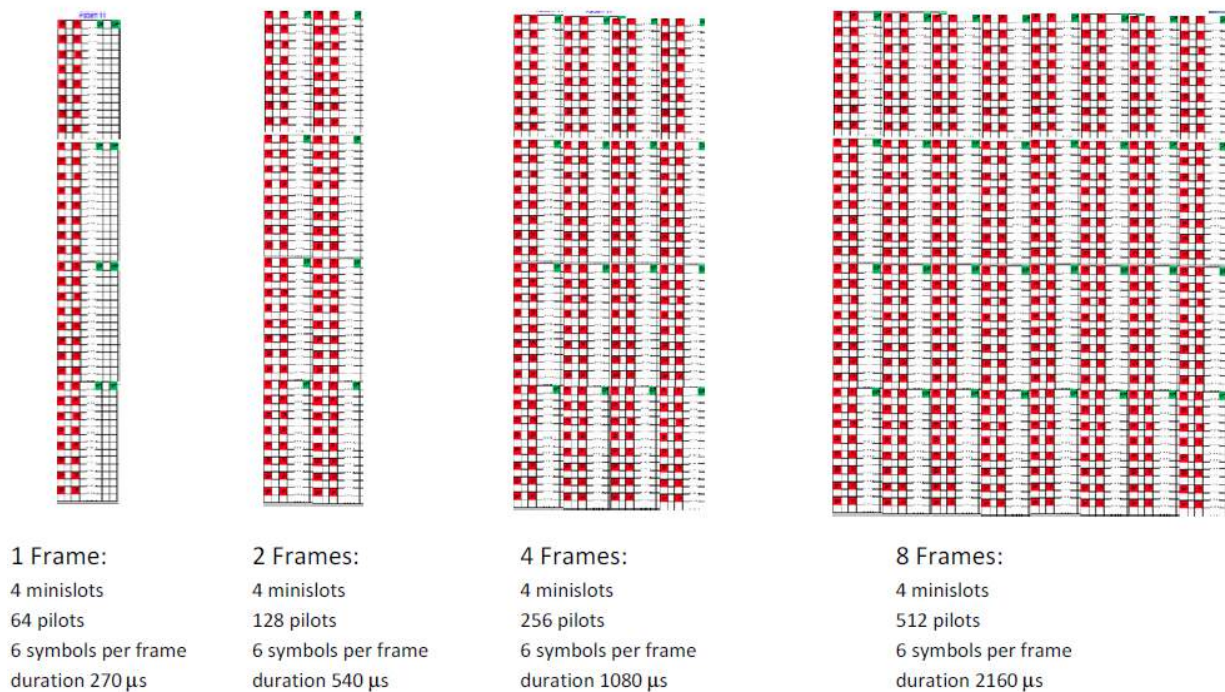


Figure 13 - Patterns of OUDP Test Burst in Time-Frequency Domain

## 8. Receiver Sensitivity analysis of the proposed detection approaches

### 8.1. Downstream Out of Band (OOB) CW pilots

The receiver sensitivity of this approach is identical to the status quo and well-known downstream leakage equipment widely used today. As such, no discussion on sensitivity is required.

### 8.2. CW Time-Division Multiple Access (TDMA)

#### Estimated sensitivity of leak detection

The sensitivity of the leak detector in the aeronautical band is a relevant parameter for FCC compliance. Technical standard § 76.605 specifies a signal leakage limit for digital signals at 17.4  $\mu$ V/m @ 3 meters. This leakage level should be measured at a bandwidth of 6 MHz for QAM, which



is why it is reasonable to assume that the energy of OFDMA leakage should be calculated in the same 6 MHz bandwidth.

Based on the above assumption, the sensitivity of OFDMA leak detector is calculated as follows:

$$S_{\text{OFDMA}} \text{ (dBmV/m)} = S_{\text{test}} \text{ (dBmV)} + \text{AF (dB/m)} + \text{N (dBc)},$$

where:

$S_{\text{test}}$  (dBmV) is sensitivity of CW test signal receiver;

AF (dB/m) is antennas factor;

N (dBc) is coefficient of recalculation level of OFDMA signal in BW= 6 MHz to the measured level of CW test signal.

The sensitivity  $S_{\text{test}}$  depends on the BW of the detected signal, receiver noise figure, and detection threshold over the noise floor. For a test signal with a duration of 1 ms (BW = 1 kHz), the maximum realized sensitivity for a detection threshold of 6 dB is around -85 dBmV.

The antennas factor (AF) for a monopole antenna at 135 MHz is around 8 dB/m.

The coefficient N is defined by the formula:

$$\text{N (dBc)} = 10 \text{ Log(M)} - \text{K (dBc)},$$

where:

M is number of subcarriers in BW = 6 MHz. M=240 for 25 kHz subcarriers spacing and M= 120 for 50 kHz subcarriers spacing;

K (dBc) is boosting gain of CW test signal.

Therefore, for 25 KHz spacing OFDMA signal (M=240) coefficient N (dBc) equals:

$$\text{N (dBc)} = 10 \text{ Log}(240) - \text{K} = 23.8 - \text{K}$$

Finally, the sensitivity of the OFDMA leak detector can be estimated as follows:

$$S_{\text{OFDMA}} \text{ (dBmV/m)} = -85 \text{ dBmV} + 8 \text{ dB/m} + (23.8 - \text{K}) \text{ dBc} = -53.2 - \text{K dBmV/m}$$

Table 3 below shows the estimated sensitivity for different boosting gains K:

**Table 3 - Estimated Sensitivity for Different Boosting Gains, K**

| Boosting gain K, dBc | Sensitivity |      |
|----------------------|-------------|------|
|                      | dBmV/m      | μV/m |
| 0                    | -53.2       | 2.19 |
| 4.7                  | -57.9       | 1.38 |
| 10                   | -63.2       | 0.69 |
| 15                   | -68.2       | 0.39 |
| 20                   | -73.2       | 0.22 |

Note, the above sensitivity estimation is for ideal conditions. In real life, in the presence of ambient noise, actual sensitivity will be less. For example, if ambient noise spectral density is anticipated to be 20 dBc above the thermal noise at the receiver, then the actual sensitivity will be also 20 dB (ten

times  $\mu\text{V/m}$ ) worse. This is why a maximum boosting of a CW test signal is preferable for robust leak detection.

### 8.3. CW Frequency-Division Multiple Access (FDMA)

#### Estimated sensitivity of leak detection

Similar to the logic of the CW TDMA approach described above, for this approach (FDMA) the energy of OFDMA leakage also should be calculated in the same 6 MHz bandwidth.

Based on this assumption, the sensitivity of OFDMA leak detector is calculated as follows:

$$S_{\text{OFDMA}}(\text{dBmV/m}) = S_{\text{test}}(\text{dBmV}) + \text{AF}(\text{dB/m}) + \text{N}(\text{dBc}),$$

where:

$S_{\text{cw\_tone}}(\text{dBmV})$  is sensitivity of CW tone receiver;

AF (dB/m) is antennas factor;

N (dBc) is coefficient of recalculation level of OFDMA signal in BW= 6 MHz to measured level of CW tone signal.

The sensitivity  $S_{\text{cw\_tone}}$  depends upon receiver noise figure, bin spacing at FFT processor, and detection threshold over noise floor. In modern leakage detectors currently utilized, the sensitivity of detection for similar CW tone signal at OFDM and Pilot/QAM modes is around -100 dBmV.

The antennas factor (AF) for a monopole antenna at 135 MHz is around 8 dB/m.

The coefficient N is defined by formulas:

$$\text{N}(\text{dBc}) = 10 \text{ Log}(M) - L(\text{dBc}),$$

where:

M is number of subcarriers in BW = 6 MHz. M=240 for 25 kHz subcarriers spacing and M= 120 for 50 kHz subcarriers spacing;

L (dBc) is relative level of CW tone signal (see Table 2).

For 25 KHz spacing OFDMA signal (M=240) coefficient N (dBc) equals to:

$$\text{N}(\text{dBc}) = 10 \text{ Log}(240) - L = 23.8 - L$$

Finally, the sensitivity of the OFDMA leak detector can be estimated as follows:

$$S_{\text{OFDMA}}(\text{dBmV/m}) = -100 \text{ dBmV} + 8 \text{ dB/m} + (23.8 - L)\text{dBc} = -68.2 - L$$

By using result for L (dBc) from Table 2, it can be seen that to provide sensitivity at around -60 dBmV/m (1  $\mu\text{V/m}$ ), a reasonable exclusion BW should be 800 kHz.

This sensitivity is sufficient for FCC compliance.

#### 8.4. OUDP Burst Test Signal

The sensitivity of the OFDMA leak detector, in the case of the OUDP approach, is calculated as follows:

$$S_{\text{OFDMA}}(\text{dBmV/m}) = S_{\text{OUDP}}(\text{dBmV}) + \text{AF} (\text{dB/m}) + \text{N} (\text{dBc}),$$

where:

$S_{\text{OUDP}}(\text{dBmV})$  is the sensitivity of the OUDP test signal receiver;

$\text{AF} (\text{dB/m})$  is the antennas factor;

$\text{N} (\text{dBc})$  is the coefficient of the recalculation level of the OFDMA signal in  $\text{BW} = 6 \text{ MHz}$  to the level of the signal at the output of the OUDP-matched filter.

The sensitivity  $S_{\text{OUDP}}$  depends on the number of pilots in the OUDP test burst, cyclic prefix duration, receiver noise figure, and detection threshold over noise floor. In a matched filter scenario, the energy of all pilots are coherently combined within time slot  $T$  of one OFDMA symbol, plus any cyclic prefix. So, the sensitivity  $S_{\text{OUDP}}$  equals the sensitivity of the detection CW burst with duration  $T$  and level boosted  $K$  times, where  $K$  is the number of pilots in the OUDP test burst:

$$S_{\text{OFDMA}}(\text{dBmV/m}) = S_{\text{CW-T}}(\text{dBmV}) - 10 \text{ Log}(K)$$

The antennas factor (AF) for a monopole antenna at 135 MHz is around 8 dB/m.

The coefficient  $N$  is defined by formula:

$$N (\text{dBc}) = 10 \text{ Log}(M),$$

where:

$M$  is number of subcarriers in  $\text{BW} = 6 \text{ MHz}$ .

For 25 KHz spacing of an OFDMA signal, the number of subcarriers  $M$  is 240, and coefficient  $N$  equals to 23.8 dBc. Thus, the sensitivity of an OFDMA leak detector for 25 KHz spacing and  $T = 45 \mu\text{s}$  (symbol 40  $\mu\text{s}$  plus cyclic prefix 5  $\mu\text{s}$ ) can be estimated as follows:

$$S_{\text{OFDMA}}(\text{dBmV/m}) = S_{\text{CW-T}} - 10 \text{ Log}(K) + 8 + 23.8 = S_{\text{CW-45}} - 10 \text{ Log}(K) + 31.8$$

The sensitivity  $S_{\text{CW-45}}$  of detection CW burst with duration 45  $\mu\text{s}$  is 13.47 dBc (in 22.2 times) worse than maximal sensitivity of detection CW burst with duration 1 ms and threshold 6 dB (see analysis sensitivity for CW burst approach). But lab tests showed that in case of OUDP the detection threshold should be increased at least to 12 dB or on + 6 dBc to prevent false alarms. Thus, the sensitivity  $S_{\text{CW-45}}$  can be estimated as follows:

$$S_{\text{CW-45}} = - 85 \text{ dBmV} + 13.47 \text{ dBc} + 6 \text{ dBc} = - 65.5 \text{ dBmV}$$

Finally, the sensitivity of the OFDMA leak detector equals:

$$S_{\text{OFDMA}}(\text{dBmV/m}) = - 65.5 - 10 \text{ Log}(K) + 31.8 = - 33.7 - 10 \text{ Log}(K)$$

As follows from the above formula, increasing the sensitivity requires increasing the number of pilots. This means increasing the number of OUDP minislots and frames, and using a pilot pattern with the maximum number of pilots (pattern 11 within the existing specifications). Table 4 shows the estimated sensitivity for pilot pattern 11 and for different combinations of minislots and frames. This sensitivity is sufficient to meet the FCC's signal leakage requirements.


**Table 4 - Sensivity for Pilot Pattern 11**

| Number of minislots / frame | Number of pilots (K) | Sensitivity |           |
|-----------------------------|----------------------|-------------|-----------|
|                             |                      | dBmV/m      | $\mu$ V/m |
| 2 /1                        | 32                   | - 48.8      | 3.6       |
| 2/2 or 4/1                  | 64                   | - 51.8      | 2.6       |
| 2/4 or 4/2                  | 128                  | -54.8       | 1.8       |
| 2/8 or 4/4                  | 256                  | -57.8       | 1.3       |
| 4/8                         | 512                  | -60.8       | 0.91      |

## 9. Using Full Band Capture to Evaluate Potential Aeronautical Band Leakage Issues

One of the features of D3.0 and consequent devices is the ability to view the downstream spectrum using full band capture (FBC). Comcast's network has 40M+ devices with FBC and this feature is widely used as part of our proactive network maintenance (PNM) program.

Prior to installing a high split device, FBC can be used to evaluate the quality of the home network by looking at ingress in the FM band from 88-108 MHz. This band is at the edge of the aeronautical band, so, high ingress will indicate the potential for leakage issues into that spectral region. The level out of the cable modem in the aeronautical band can be more than 2000 times higher than the downstream level into the cable modem. Below are two examples of FBC screen captures showing ingress. In Figure 14, the noise floor in the FM band is very low, indicating a home with high network quality, that is a good candidate for high split service with potentially little leakage.



**Figure 14 - FBC Capture Showing No Ingress in FM Band**

Figure 15 shows high ingress and the potential for leakage if converted to a high split and transmitting in the aeronautical band. In this instance, remediation can be performed prior to providing high split service, to minimize leakage issues in the FM band.

Full Band Capture Showing High Ingress in FM Band




Figure 15 - FBC Showing Ingress in the FM Band

## 10. Probability of Intercept: Moving Vehicle Analysis

If upstream test signals used for leakage detection are bursty in nature, a careful examination needs to be conducted to assess how many bursts can be detected by a moving vehicle (in a one second time interval), equipped with a leakage detector. The probability of intercept (POI) is the probability of capturing a transmitted signal based on Tx/Rx timing and should be 100% of the time. With the burst signal approaches, including burst CW or OUDP, the transmit duration and cycle time can be configured such that each CM transmits approximately 2x per second. The real time FFT within the leakage detector is continuously on, so with certainty (timing-wise), signal detection will occur if the amplitude of the signal at the detector is greater than the sensitivity of the detector. See Figure 16.

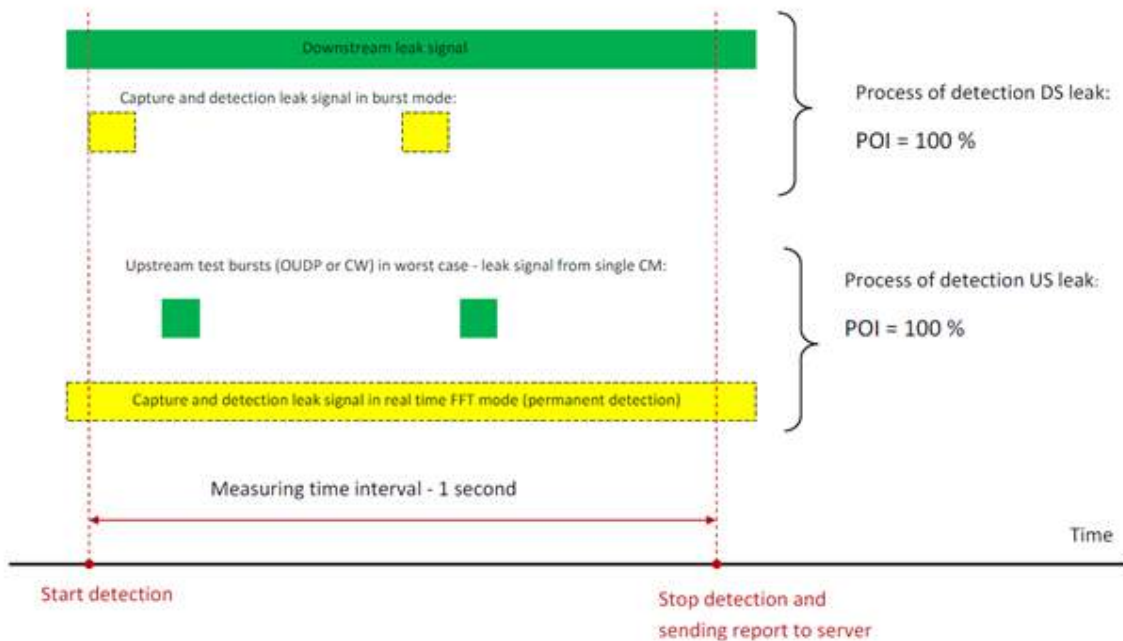
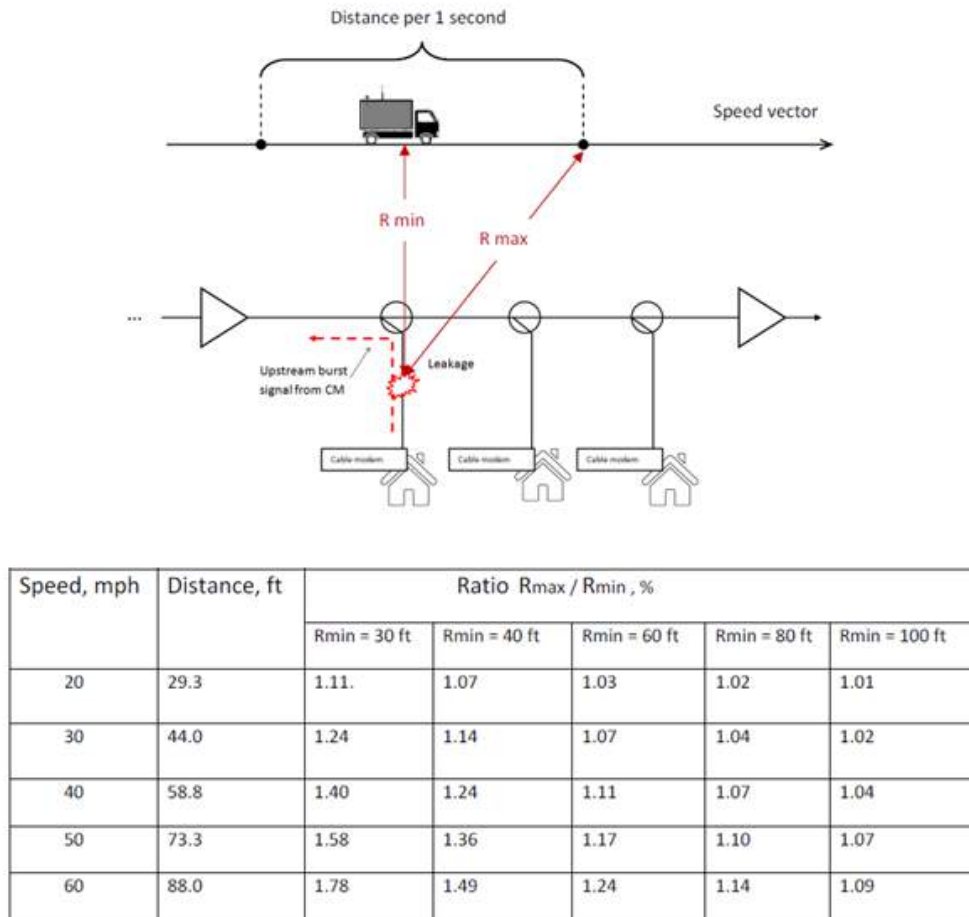


Figure 16 - Probability of Intercept (POI), Detection of DS and US leaks from moving vehicle


Figure 17 show the results of a model that was built to compare the ratio of the minimum to maximum distance a vehicle could have travelled before receiving a burst signal. The ratio allows an estimate of the change of a leaking signal's level, at a vehicular Rx point, during the 1 second measuring session. The maximum ratio of 1.78 applies to a max speed of 60 mph and a minimum distance (Rmin) of 30 feet. Definitely not a realistic scenario, but even in this case, the field strength variation will not be so big. For a typical scenario, of a vehicular speed of 30 mph and an Rmin of 60 ft, the ratio will be only 1.07. It means the same as changing 1.07 in  $\mu\text{V}/\text{m}$  or 0.58 dB in dBmV/m, a trivial difference.



**Figure 17 - Distance covered in 1 second versus vehicle speed and ratio of the minimum distance to the maximum distance a vehicle could have travelled**

Using real time signal processing for upstream leakage detection allows the MSO to provide a 100% probability of burst leak capture (POI). The variation of leak level(s) at the moving vehicle will be within one dB and therefore can be ignored. With high-split cable modems following a gradual rollout, the total time needed to have all cable modems in a service group (SG) to burst leakage detection signals is lowered, allowing the 1.6 MHz BW used for the OUDP burst to transmit upstream data part of the time, as illustrated in Figure 18 - Leakage Detection OUDP Test Bursts and Data Transmission for a SG below. This technique of maximizing the use of the spectrum for data, while also ensuring the reliability of leakage compliance, helps meet the goals of moving to the 204 MHz split for both capacity and 1 Gbps

upstream performance. This scheduling capability can be easily facilitated with the virtual CMTS (vCMTS) architecture; other CMTS architectures can provide similar capabilities.



**Figure 18 - Leakage Detection OUDP Test Bursts and Data Transmission for a SG**

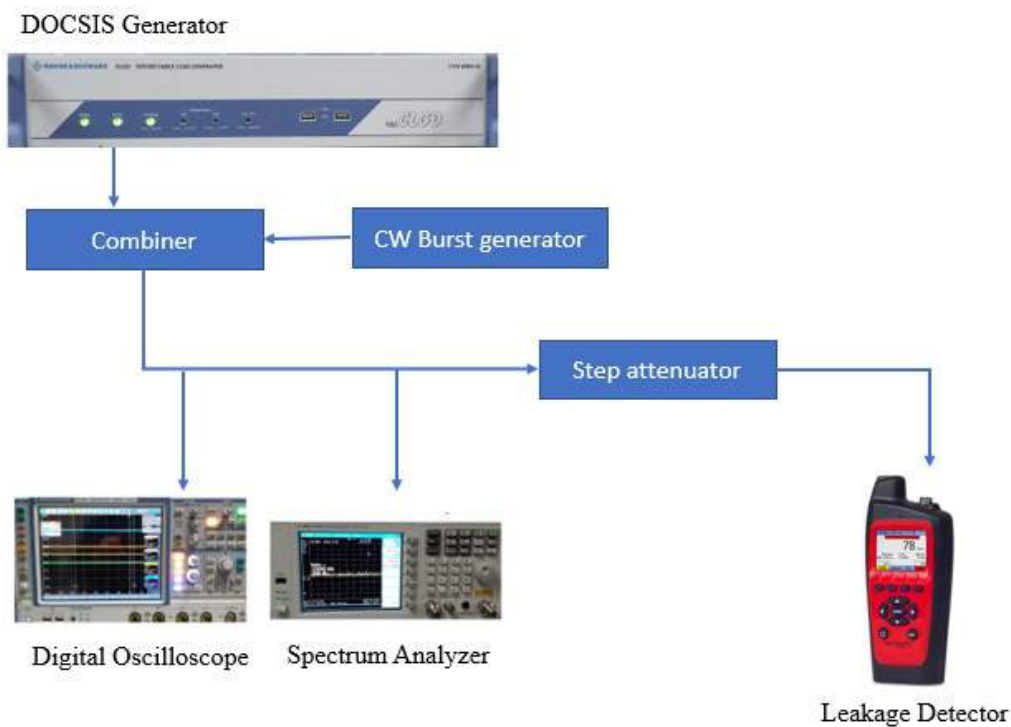
## 11. Digital Leakage Detection Test Results

Testing was performed on the various methods described within this paper to validate the concepts and determine minimum sensitivities for each.

### 11.1. CW Time-Division Multiple Access (TDMA)

To prove out the proof of concept for CW TDMA leakage detection, the test setup in Figure 19 was utilized. A DOCSIS generator supplied the OFDMA channel, and to create the exclusion zone within which the CW Burst carriers were placed. The test signal was directly connected to the leakage detector.

The CW test burst in the time and frequency domain are shown in Figure 20 below.



**Figure 19 - CW Leakage Detection Test Setup**



**Figure 20 - CW test burst in time domain and spectrum of CW bursts**

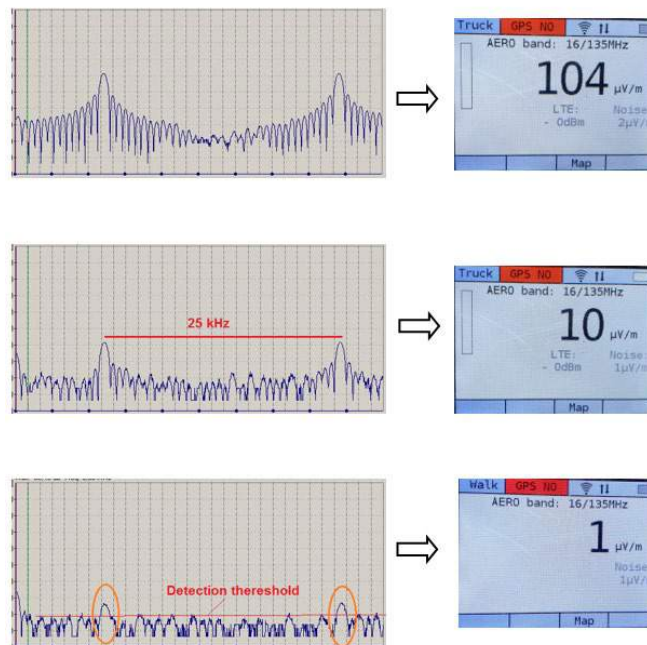


The test setup occupied a 4-subcarrier exclusion zone within the OFDMA, for an exclusion bandwidth of 100 kHz. The spectrum of the CW carriers within this exclusion bandwidth is shown in Figure 21.



**Figure 21 - Spectrum of CW test bursts within exclusion BW of OFDMA signal**


The FFT detection results and the corresponding leakage detection display results are shown in Figure 22. Boosting gain was  $K=+10\text{dBc}$ . In the top image, the detection result was  $104\mu\text{V/m}$ . The CW carrier spacing is 25kHz. The signals were attenuated to a detected level of  $10\mu\text{V/m}$ , and further attenuated to  $1\mu\text{V/m}$ , which is just above the detection noise floor.



**Figure 22 - FFT spectrums and results of leak detection for boosting gain  $K=+10\text{ dBc}$**



An example of how the detector was configured is shown in Figure 25, importantly showing that with proper configuration, the detector is able to detect all the different tones simultaneously, with no change in detector setting required. An example of the detection test result of the minimum detected signal level is shown below.



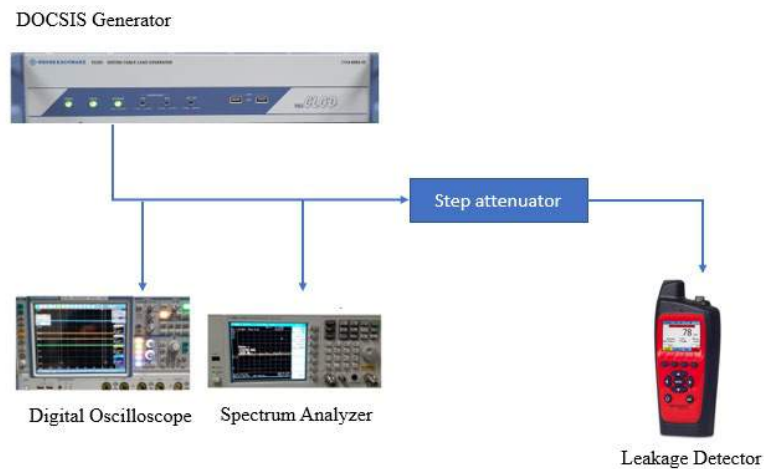
**Figure 25 - Leakage Detector settings for detection of CW tones and minimum detected signal**

Sensitivity of detection OFDMA leakage in the aeronautical band is approximately 1...2mV/m and can be provided by selection of an exclusion bandwidth of 800 kHz, which is a reasonable capacity loss while still meeting the goals of the high split. In this case, the CW tones level will be approximately 6 to 8 dB below the level of OFDMA subcarrier, or approximately - 30 dBc below energy of OFDMA signal in BW = 6 MHz. It's a very similar scenario of the status quo detection of pilots between QAMs in the downstream, with well-known results.


### 11.3. OUDP Burst Test Signal

To prove out the proof of concept of the OUDP Burst Test Signal detector, a method was needed to simulate the OUDP burst. Again, a DOCSIS generator supplied the test signal, which, after an attenuator, was directly connected to the leakage detector. The test setup is shown in Figure 26.

The test pattern supplied by the DOCSIS generator is shown in Figure 27.


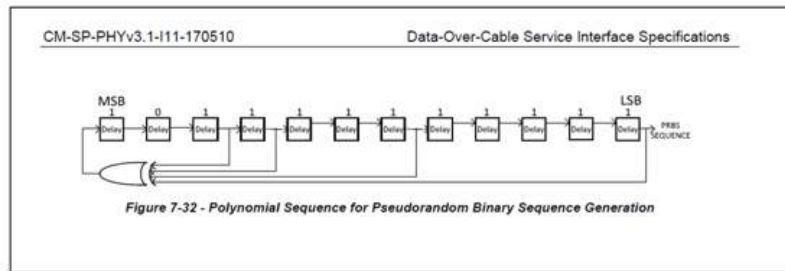


**Figure 26 - Test bench block diagram**



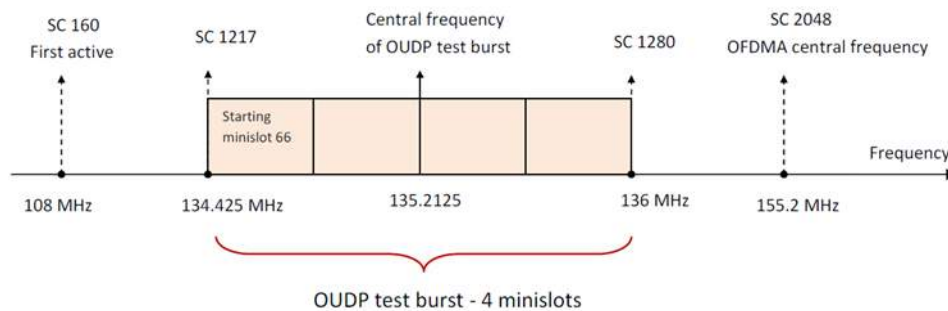
**Figure 27 - Pattern of the OUDP test burst in time – frequency domain**

The modulation of the OUDP test burst is shown in Figure 28.



**Figure 28 - Pilot's BPSK modulation in OUDP test burst in Figure 27 [1]**

shows how the OUDP burst signal appears in the frequency domain, including central frequency, minislots, and subcarriers.



**Figure 29 - Allocation of OUDP test burst in frequency domain of aeronautical band**

Figure 30 - Arbitrary waveform generation settings at DLCG for simulation OUDP test burst shows the settings that were used in the configuration of the DOCSIS generator to simulate the OUDP burst signal.



**DOCSIS Cable Load Generator**  
Upstream Arbitrary Waveform Generation & Transmission

Mode: **UPSTREAM**  
D3.0 Mode: **OFF**  
Cmd Mode: **WEB**

Release Control

Home  
DOCSIS 3.0  
DOCSIS 3.1  
Downstream ARB  
**Upstream**  
Impairments  
Network  
User Files  
Licences  
Error Queue  
Help  
About

Preset  
Generate  
Undo

**Generation** | Transmission | Upstream Impairments

Type: OFDMA  
Burst Type: Data Mode

| Parameter                  | Value                  | Parameter | Value |
|----------------------------|------------------------|-----------|-------|
| FFT Size                   | 4K                     |           |       |
| Encompassed Spectrum       | 160                    |           |       |
| First Active Subcarrier    | 3900                   |           |       |
| Last Active Subcarrier     | 3900                   |           |       |
| Cyclic Prefix (us)         | 5.0                    |           |       |
| Window Rolloff Period (us) | 0.9375                 |           |       |
| Symbols Per Frame K        | 8                      |           |       |
| Pilot Pattern              | 11                     |           |       |
| Modulation Order           | QPSK                   |           |       |
| Scrambler                  | ON                     |           |       |
| Scrambler Seed             | minislotAdvancedButton |           |       |

Exclusion Bands

| Start | Width |
|-------|-------|
| 0     | 0     |
| 0     | 0     |
| 0     | 0     |

Data Mode

Number of Frames: 1  
User Starting Minislot in Frame: 66  
User Minislots Per Frame: 4  
Data Filename:   
Fill Remaining Minislots: OFF

- Key settings

**Figure 30 - Arbitrary waveform generation settings at DCLG for simulation OUDP test burst**

The configuration of the RF transmission of the DOCSIS generator was performed as in Figure 31 - Arbitrary waveform transmission settings at DCLG for simulation OUDP test burst.

**DOCSIS Cable Load Generator**  
Upstream Arbitrary Waveform Generation & Transmission

Mode: **UPSTREAM**  
D3.0 Mode: **OFF**  
Cmd Mode: **WEB**

Release Control

Home  
DOCSIS 3.0  
DOCSIS 3.1  
Downstream ARB  
**Upstream**  
Impairments  
Network  
User Files  
Licences  
Error Queue  
Help  
About

Preset  
Apply  
Undo  
Trigger Configuration  
Trigger

**Generation** | **Transmission** | Upstream Impairments

Impairment Destination Channel: OFF

| Channel | Transmit | Power (dBmV) | Frequency (MHz) | Mode       | Output Delay (us) | Inter burst Gap (us) | ARB File                      |
|---------|----------|--------------|-----------------|------------|-------------------|----------------------|-------------------------------|
| 1       | ON       | 30.0         | 155.2           | Triggered  | 0                 | 5000                 | Arcom POC 4k_11_Frame x 4.wav |
| 2       | OFF      | 30.0         | 80              | Continuous |                   |                      |                               |
| 3       | OFF      | 30.0         | 120             | Continuous |                   |                      |                               |
| 4       | OFF      | 30.0         | 160             | Continuous |                   |                      |                               |
| 5       | OFF      | 30.0         | 160             |            |                   |                      |                               |
| 6       | OFF      | 30.0         | 160             |            |                   |                      |                               |
| 7       | OFF      | 30.0         | 160             |            |                   |                      |                               |
| 8       | OFF      | 30.0         | 160             |            |                   |                      |                               |
| 9       | OFF      | 30.0         | 160             |            |                   |                      |                               |
| 10      | OFF      | 30.0         | 160             |            |                   |                      |                               |
| 11      | OFF      | 30.0         | 160             |            |                   |                      |                               |

Trigger Configuration

Trigger Source: INTERNAL  
Trigger in Edge Sensitivity: RISE  
Trigger out Start Delay (us): 0  
Trigger out Active Polarity: HIGH

Apply Cancel

- Key settings

**Figure 31 - Arbitrary waveform transmission settings at DCLG for simulation OUDP test burst**

The simulated OUDP burst signal in the time and frequency domain is shown Figure 32 .

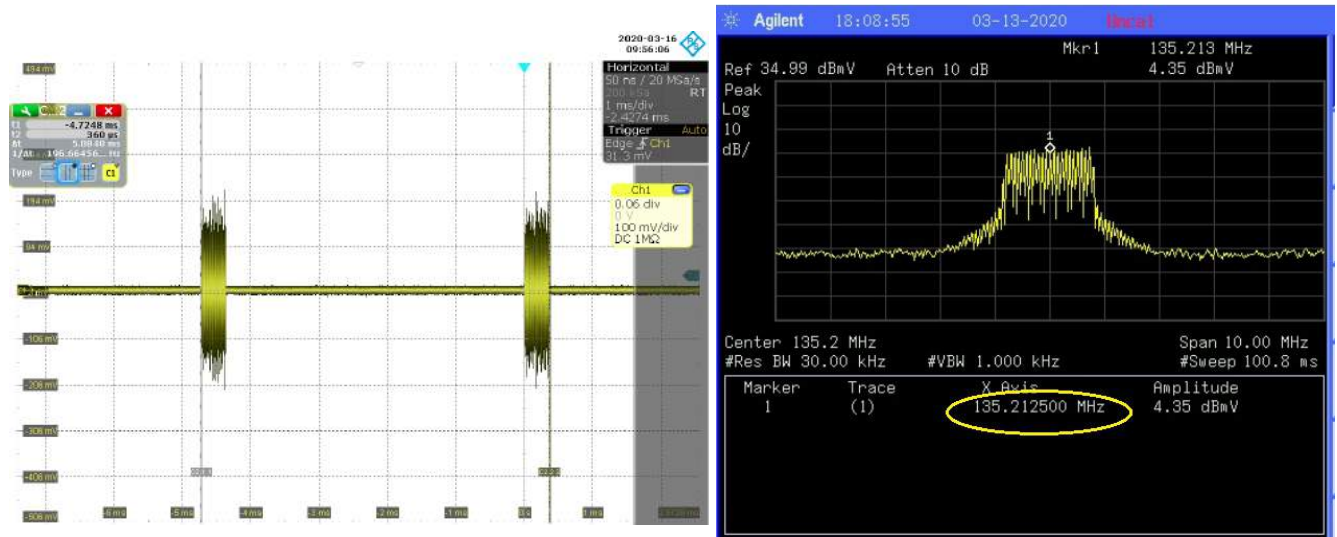


Figure 32 - OUDP test burst in time and frequency

The response of the matched filter in the time domain is shown below. The peak of the cross-correlation function represents the detected leak magnitude. The envelope of the OUDP burst signal is also visible.


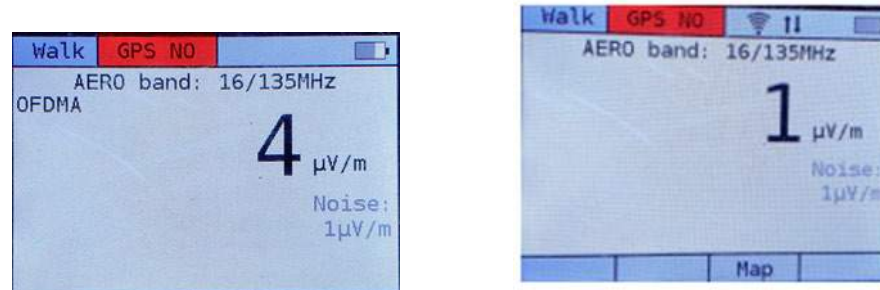


Figure 33 - Response of matched filter for OUDP test burst with pilot pattern 11.

In Figure 34 , the minimal detected leak level (sensitivity) for 4 minislots, with one frame in the conducted test, was 4 $\mu$ V/m as shown on the left. In order to investigate improved sensitivity, the test setup was re-run using 4 minislots and eight frames. Sensitivity was improved to 1 $\mu$ V/m, as shown on the right.



**Figure 34 - Results of implementation of UUDP burst detector into field meter**

This approach, of matched filtering detection of the OUDP pilot pattern 11, is sufficient to meet FCC requirements. The estimated sensitivity detection of the OUDP test bursts burst is approximately 4  $\mu$ V/m for 4 minislots (BW = 1.6 MHz) with one frame, and approximately 1  $\mu$ V/m for 4 minislots (BW = 1.6 MHz) with 8 frames. Therefore, operators will have an opportunity to select the optimal BW of OUDP test bursts for leak detection, depending upon desired sensitivity.

## 12. Conclusions

As cable systems evolve and the need for more bandwidth continues, high split systems are now needed; FDX systems that will be used in the near future and are presently being developed.

Current leakage methods are well established, and the need to build off these legacy techniques, while evolving to support new technologies and spectral designations is both necessary and required by federal regulations.

These new systems, and the inverted signal levels that high split systems create, pose unique challenges to leakage measurements and FCC compliance that have not been considered in the 60+ years of cable and broadband evolution. As discussed in this paper, there are several options to move forward with leakage measurements. In addition, operators need to evaluate FCC requirements based on signal type and transmission level in the aeronautical band.

The initial stance to detect leakage in a high split system is typically to continue to use and measure leakage tones in the downstream direction. A leak is a leak, right? However, the differences in signal levels from leakage tones in the downstream, vs actual upstream signals generated at the cable modem, are difficult to correlate. Complex GPS and system mapping algorithms are needed. Acceptance by the FCC for using calculated leakage values and not measuring in the aeronautical band will be difficult. All of the other proposed methods have shown, from a detection sensitivity perspective, to be sufficient for FCC compliance purposes.

Using CW tones -- either continuous or burst generated in the cable modems -- seems like the next logical solution. Cable modems are not designed to create CW tones, and they are designed specifically not to continuously transmit signals in the upstream band. If implementing one of the CW approaches is



desired, certainly the logistics required for the CM software updates would be formidable, and the cost of doing so would need to be analyzed. If the CW-TDMA approach is implemented, additional complexity involving CMTS capabilities would be required, in the form of controlled CM burst timing.

If the CW-FDMA approach is of interest, there exists the additional complication of tracking frequency designations for each CM in a node – a challenging task, made more so by the common practice of node splitting. One additional benefit of the CW-FDMA approach is that even if only the Tx frequency of one CM in a node is known, it is possible to precisely measure the frequency of the two detected CW leakage tones, and therefore exactly resolve which CM supplied the detected leak. This could result in multiple operational benefits.

In a High Split system, it is advantageous to use OFDMA blocks for as much of the spectrum as possible, for maximum upstream throughput. Measuring the OFDMA signal from normal traffic would be the ideal solution, but this is not feasible. DOCSIS 3.1 provides for an OUDP test burst to be generated, and this is the ideal signal to use for leakage measurements. It is configurable in both frequency and duration and can be set to the same transmit level as the cable modem. It provides the tools necessary to manage both the capacity and the leakage compliance, with flexibility for the operator. The pilot pattern can be chosen and configured such that the sensitivity of this signal at the detector is equivalent to the CW approaches currently used in the downstream.

Table 6 shows a summary of the solutions evaluated in this paper. Take note that the OUDP approach is already supported by both the CM firmware and DOCSIS specifications. This approach is also supported in FDX, and the frequency of the OUDP test burst can be configured higher in frequency as the upstream band is extended in FDX and other topologies.

**Table 5 - Summary of Leakage Detection Methods**

|                                                                             | Leak Detection Method |         |         |          |
|-----------------------------------------------------------------------------|-----------------------|---------|---------|----------|
|                                                                             | OUDP                  | CW-TDMA | CW-FDMA | DS Pilot |
| Detected signal level correlated with actual aeronautical band signal level | yes                   | yes     | yes     | no       |
| Detection sensitivity sufficient for FCC compliance                         | yes                   | yes     | yes     | yes      |
| Realizable within current DOCSIS specification                              | yes                   | no      | no      | no       |
| Realizable with current firmware in all High Split cable modems             | yes                   | no      | no      | yes      |
| Requires permanent orchestration from the CMTS                              | yes                   | yes     | no      | no       |
| Requires some dedicated US Bandwidth                                        | yes                   | yes     | yes     | no       |

With the High Split leakage options evaluated in this paper, operators can derive a roadmap for continued leakage monitoring that not only allows for continued regulatory compliance, but also consistent plant monitoring and maintenance for optimal system performance.

## Abbreviations

|      |                              |
|------|------------------------------|
| AF   | antenna factor               |
| AP   | access point                 |
| bps  | bits per second              |
| BW   | bandwidth                    |
| CATV | community antenna television |
| CFR  | code of federal regulations  |
| CLI  | cumulative leakage index     |

|         |                                                                                                          |
|---------|----------------------------------------------------------------------------------------------------------|
| CM      | cable modem                                                                                              |
| CMTS    | cable modem termination system                                                                           |
| COVID   | corona virus disease                                                                                     |
| CPE     | customer Premises equipment                                                                              |
| CW      | continuous wave                                                                                          |
| DAA     | distributed access architecture                                                                          |
| dBc     | decibel from Carrier                                                                                     |
| dBmV    | decibel Millivolt                                                                                        |
| DOCSIS  | data over cable service interface specification                                                          |
| DSG     | DOCSIS set-top gateway                                                                                   |
| DS      | downstream                                                                                               |
| ESD     | extended spectrum DOCSIS                                                                                 |
| FCC     | Federal Communications Commission                                                                        |
| FEC     | forward error correction                                                                                 |
| FDC     | forward data carrier                                                                                     |
| GPS     | global positioning system                                                                                |
| HFC     | hybrid fiber-coax                                                                                        |
| f       | frequency                                                                                                |
| FBC     | full band capture                                                                                        |
| FDX     | full duplex DOCSIS                                                                                       |
| FFT     | fast fourier transform                                                                                   |
| FM      | frequency modulation                                                                                     |
| Gbps    | gigabit per second                                                                                       |
| HD      | high definition                                                                                          |
| Hz      | hertz                                                                                                    |
| IOT     | internet of things                                                                                       |
| ISBE    | International Society of Broadband Experts                                                               |
| K       | maximum number of cable modems on a node                                                                 |
| K       | symbols per frame                                                                                        |
| K (dBc) | boosting gain of CW test signal                                                                          |
| KHz     | kilohertz                                                                                                |
| L       | relative level of CW tone signal                                                                         |
| m       | meter                                                                                                    |
| M       | number of excluded sub carriers                                                                          |
| M       | number of subcarriers in 6 MHz                                                                           |
| MAC-PHY | media access control channel physical layer                                                              |
| Mbps    | Megabit per second                                                                                       |
| MHz     | megahertz                                                                                                |
| mS      | millisecond                                                                                              |
| MSO     | multiple system operator                                                                                 |
| MVPD    | multichannel video programming distributor                                                               |
| N (dBc) | coefficient of recalculation level of OFDMA signal in BW= 6 MHz to the measured level of CW test signal. |
| NCTA    | National Cable Television Association                                                                    |
| OOB     | out of Band                                                                                              |
| OFDM    | orthogonal frequency division multiplexing                                                               |
| OFDMA   | orthogonal frequency division multiple access                                                            |
| OTT     | over-the-top                                                                                             |

|                |                                                |
|----------------|------------------------------------------------|
| OUDP           | OFDMA upstream data profile                    |
| POI            | probability of intercept                       |
| QAM            | quadrature amplitude modulation                |
| RF             | radio frequency                                |
| RPD            | remote physical device                         |
| R-PHY          | remote physical layer                          |
| RX             | receive                                        |
| SC-QAM         | single carrier quadrature amplitude modulation |
| SCTE           | Society of Cable Telecommunications Engineers  |
| $S_{cw\_tone}$ | sensitivity of CW tone receiver                |
| $S_{OFDMA}$    | sensitivity of OFDMA leak detector             |
| $S_{test}$     | sensitivity of CW test receiver                |
| STB            | set-top box                                    |
| SCW            | sensitivity detection of CW burst              |
| $S_{cw\_tone}$ | sensitivity detection of CW tone               |
| SG             | service group                                  |
| STD            | standard                                       |
| TCP            | total composite power                          |
| TDMA           | time division multiple access                  |
| uV/m           | microvolt per meter                            |
| uS             | microsecond                                    |
| US             | upstream                                       |
| Wi-Fi          | wireless fidelity                              |

## Bibliography & References

- 1) [PHYv3\_1] Physical Layer Specification, CM-SP-PHYv3.1-I17-190917, September 09, 2019, Cable Television Laboratories, Inc
- 2) In signal processing, a matched filter is obtained by correlating a known delayed signal, or *template*, with an unknown signal to detect the presence of the template in the unknown signal.
- 3) *Another Look at Signal Leakage, The Need to Monitor at Low and High Frequencies*; Ron Hranac, Greg Tresness, SCTE EXPO '12
- 4) Code of Federal Regulations, Title 47, Part 76 MULTICHANNEL VIDEO AND CABLE TELEVISION SERVICE

# Critical Facility Cooling Energy Optimization

A Technical Paper prepared for SCTE•ISBE by

**Thomas Hurley**

Principal Mechanical Engineer  
Comcast Cable Corporation  
55 Executive Drive, Hudson NH 03563  
603-481-0909  
thomas\_hurley@comcast.com

**John Dolan**

Senior Guideline Specialist  
Rogers Communications Inc.  
8200 Dixie Road, Brampton ON CA L6T 0C1  
519-852-5666  
john.dolan@rci.rogers.com

**Arnold Murphy**

President  
Strategic Clean Technology Strategic Clean Technology Inc. (SCTi)  
3476 Galetta Road, Amprior, ON K7S 3G7  
613-558-4415  
a.murphy@sct-inc

**Mike Glaser**

Critical Facilities Engineer IV  
Cox Communications Inc.  
6305B Peachtree Dunwoody Rd. Atlanta, GA 30328  
404-427-5302  
mike.glaser@cox.com

**John Teague**

Director Strategic Solutions  
Worldwide Environmental Services  
430 Virginia Drive, Fort Washington, PA 19034  
215-619-0980  
john.teague@wes.net

**Ken Nickel**

Executive Vice President  
Quest Controls, Inc.  
Corporate address: 208 9th Street Dr. West, Palmetto FL 34221  
775-409-4312  
knickel@questcontrols.com

# Table of Contents

| Title                                             | Page Number |
|---------------------------------------------------|-------------|
| 1. Introduction.....                              | 3           |
| 2. Historical Perspective.....                    | 3           |
| 3. Case studies and results.....                  | 6           |
| 3.1. HVAC Set Point Adjustments.....              | 6           |
| 3.1.1. Set Points And Sensors .....               | 6           |
| 3.1.2. ASHRAE Standards And Operating Ranges..... | 8           |
| 3.1.3. Capacity Impact .....                      | 9           |
| 3.1.4. Set Point Modifications And Results .....  | 10          |
| 3.2. Blanking Panel Installation .....            | 11          |
| 3.2.1. Types Of Blanking Panels .....             | 12          |
| 3.2.2. Airflow Efficiency.....                    | 13          |
| 3.2.3. Blanking Panel Case Study Results .....    | 14          |
| 3.3. HVAC Economizer Options.....                 | 15          |
| 3.1.1 Economizer Operations .....                 | 16          |
| 4. Conclusion .....                               | 19          |
| Abbreviations.....                                | 20          |
| Bibliography & References .....                   | 20          |

## List of Figures

| Title                                                                                                                         | Page Number |
|-------------------------------------------------------------------------------------------------------------------------------|-------------|
| Figure 1 - Cumulative Industry Savings Through 2018 for STBs and SNE Efficiency Compliance (D+R international/ NCTA.com)..... | 4           |
| Figure 2 - Cable Operators Power Use Distribution (SCTE) .....                                                                | 5           |
| Figure 3 - Data Center Power Use Distribution (Shehabi, 2016).....                                                            | 6           |
| Figure 4 - Power Monitoring Dashboard Example From Foreseer®.....                                                             | 7           |
| Figure 5 - HVAC Set Point Monitoring And Operational Display From Foreseer .....                                              | 8           |
| Figure 6 - 2015 Recommended And Allowable Envelopes For Air-Cooled Equipment .....                                            | 9           |
| Figure 7 - Vertiv Application And Performance Comparison .....                                                                | 10          |
| Figure 8 - Vertiv Performance And Set Point Display With Function Key .....                                                   | 10          |
| Figure 9 - Open RUs On Left, PlenaForm Systems Blanking Panels Installed On Right.....                                        | 12          |
| Figure 10 - Upsite Technologies 2 RU White Molded Blanking Panel .....                                                        | 12          |
| Figure 11 - Polargy PolarFlex Sheet Blanking Panel .....                                                                      | 13          |
| Figure 12 - Visual Display of Airflow Anomalies.....                                                                          | 14          |
| Figure 13 - Hours With Ideal Conditions For Air-side Economization (DOE And The Green Grid).....                              | 16          |
| Figure 14 - Aeon Dashboard With Operational Performance Of Economizer .....                                                   | 17          |
| Figure 15 - Energy Use Comparison With And Without Economizer Activated.....                                                  | 18          |

## List of Tables

| Title                                                                                | Page Number |
|--------------------------------------------------------------------------------------|-------------|
| Table 1 - Site Energy Reduction Percentages After Set Point Increase .....           | 11          |
| Table 2 - Site Energy Reduction Percentages After Blanking Panels Installation ..... | 15          |
| Table 3 - 5 Month Projected Seasonal Energy Savings From Economizer Engagement.....  | 19          |

## 1. Introduction


With the current state of the world in turmoil due to the COVID-19 virus, there has never been a time in history when cable, internet and communication networks have been so heavily stressed. Corporations, businesses, schools, healthcare and families have been thrust into new ways of meeting, learning, socializing and transacting business -- with almost all of them flowing through cable and Internet providers. With this massive increase in traffic on all of the broadband networks comes an increased use of electricity to power all of the components needed to keep the critical infrastructure (CI) up and running. This increased consumption of energy to deliver internet and cable television services is being offset by continued conservation efforts across the industry. This paper will identify how Comcast is contributing to those efforts with three case study results which show substantial savings potential if deployed across the cable operator industry.

The production spaces, where the servers, modems, switches and routers are housed, are called headends, which divide their energy usage into 2 main categories. The first is IT production load (modems, switches etc.) and second is the cooling load, with both expressed in kilowatts (kW). This paper will look at our efforts to reduce energy consumption from the cooling load at the headends. The pilot initiatives include set point adjustments on air conditioners, blanking panels in racks, and economizer options on air conditioners (also known as “free cooling”). Although these strategies have been publicly promoted in the past by trade groups, engineering firms and vendors, this paper will provide specific applications and verifiable results which have far exceeded the expectations of this author.

These programs will demonstrate saving opportunities primarily associated with computer room air conditioner (CRAC) cooling energy, ranging from **12% to 50%** without any negative impact to operations or reliability. These significant savings will be shown to measurably reduce operating expenses, carbon emissions and preserve capital with payback periods from less than 1 year to 3.5 years.

## 2. Historical Perspective

The issue of energy efficiency has long been a Comcast initiative, as evidenced by the partnership with the SCTE (Society of Cable Telecommunications Engineers), CableLabs and corporate commitments to building a more resource-efficient company, including the retention of a Chief Sustainability Officer in 2016. Comcast was a signatory to industrywide voluntary agreements in 2012 for energy efficiency improvement to set top boxes (STB) and again in 2015 for energy efficiency of small network equipment (SNE). The substantial savings from the STB initiatives across the entire cable operator companies are expressed in Figure 1 . Additional information can be found at [www.energy-efficiency.us](http://www.energy-efficiency.us) and [www.cablelabs.com](http://www.cablelabs.com).



Source: D+R International

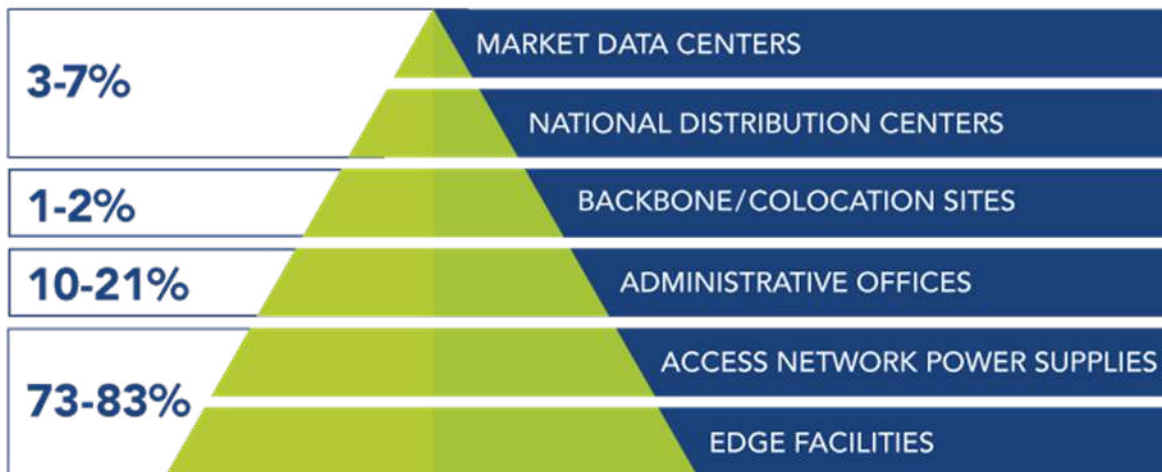
[ncta.com](http://ncta.com) | [@NCTAitv](https://twitter.com/NCTAitv)

**Figure 1 - Cumulative Industry Savings Through 2018 for STBs and SNE Efficiency Compliance (D+R international/ NCTA.com)**

In its continued drive for efficiency and environmental impact reductions, Comcast agreed to participate in the SCTE Energy 2020 program, developed in 2014 as a multi-year program. This objective covered all aspects of the industry product delivery systems with three main goals to reduce energy use. These Energy 2020 goals include reducing power consumption by 20% per unit, reduce energy cost by 25% per unit and reduce grid dependency by 5%, by 2020. With the adoption of these goals Comcast and the many other members of SCTE, including vendors and contractors, joined forces with their respective Sustainability Offices and embarked on aggressive energy reduction strategies, including the case studies referenced herein, to provide measured and verified data to various applications for cable operators.

The Energy 2020 program initiatives took a hard look at the breakdown of the major areas of energy consumption in the cable operators' system, and it is abundantly clear that outside-plant access network, and specifically power supplies and edge facilities (aka hubs & headends), represent the major energy user groups for cable and internet, as shown in Figure 2.

## CABLE OPERATOR POWER CONSUMPTION PYRAMID




Source: SCTE analysis of available Energy 2020 participating MSO data  
 © 2016 Society of Cable Telecommunications Engineers, Inc. All rights reserved.  
 scte.org • isbe.org

**Figure 2 - Cable Operators Power Use Distribution (SCTE)**

In the single use edge facilities referenced above (also known as Headends), the power breakdown consists of the production energy load (servers, modems, routers & switches), typically utilizing 45% to 55% of overall space electricity, and the Heating Ventilation and Air Conditioning (HVAC) load, using 40% to 45%. Miscellaneous office and lighting load use 3% to 5%. These loads are expressed in Kilowatts (kW) for point-in-time measurements, or, more accurately as kilowatt Hours (kWh) to capture the amount of Watts consumed by a device over 1 hour, that is, the energy consumption. The cost varies widely across the US, from \$.06 to \$.023 per kWh. California and the northeast states are in the highest electricity cost groups. It is also more likely that the higher priced areas will have more aggressive utility incentives, and the three case studies included below have been eligible for incentives either individually or as part of site upgrades. See the following link for state incentive opportunities: <https://www.dsireusa.org/> with a focus on the HVAC systems. In most incentive programs the measures addressed by this document typically are included under the Commercial & Industrial programs (C&I). In this paper we will be looking only at the reduction of kWh from HVAC cooling power usage.

Many papers have been written on the subject of energy conservation measures (ECMs) for data centers. In particular, this very similar industry attracted worldwide attention in 2014 for reports that it was consuming nearly two percent of the total energy consumption (Arman Shehabi, 2016) with projected staggering growth. This growth did not materialize, thanks to the many efficiency efforts of manufacturers, engineers and end users. Data centers require many of the same infrastructure equipment as cable headends, and efficiency applications discussed in this work are very much applicable, with just as promising a savings potential. As indicated in Figure 3, a data center contains many more servers and storage elements for business applications (e.g. accounting, human resources, marketing, purchasing, legal), while the cable headends contain modems, switches, fiber terminations, router and other broadband equipment. Although slightly different in rack contents, the energy use percentages compare fairly closely.





**Figure 3 - Data Center Power Use Distribution (Shehabi, 2016)**

A 2017 SCTE Expo technical paper titled “ECM Recommendations for Cable Edge Facilities” (*Dan Marut Senior Manager of Sustainability, 2017*), detailed similarly combined ECMs that had been deployed, and the savings aggregated. This author is in full support of the recommendations from that paper and intends to show specific individual element savings.

### **3. Case studies and results**

#### **3.1. HVAC Set Point Adjustments**

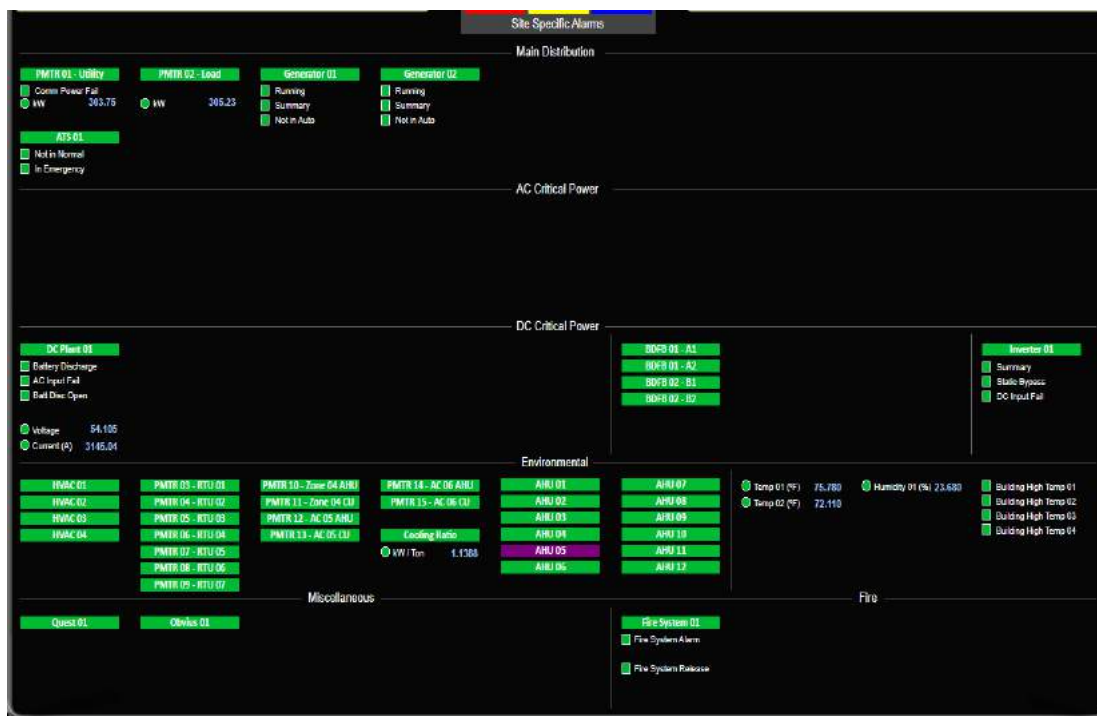
In this case study, 10 sites within the greater Chicago metro area were identified as candidates to raise the current HVAC set points in small (600 square feet) to medium (1,200 square feet) size power rooms. These sites were selected to minimize any adverse or customer-impacting events. Simultaneous with this effort was the installation of lockable and programmable thermostats used to control the CRAC devices. The installation of these secure thermostats ensured that they could not be adjusted by occupants. The addition of these devices served as qualification for a significant Commonwealth Edison (Com Ed) utility incentive. At the outset of the project, we estimated a potential energy savings of 3% to 5% for the cooling energy used, that would be attributable to the raising of return air temperature set points on the air conditioners from 68°F (Fahrenheit) to 75°F.

##### **3.1.1. Set Points And Sensors**

Set points are the numeric control positions used to operate HVAC and many other types of equipment. The most common example is a thermostat which controls when an air conditioner or heater comes on or shuts off. The difference between the settings for when a unit cycles on and off is referred to the set point deadband. When conditions are within the deadband range, no action (to provide heat or cooling) is needed by the device. In headends, the air conditioners have historically run 24/7 365 days a year to maintain a consistent operating temperature in the room. This continuous duty requires greater reliability than a seasonal device, which may run 3-6 months a year.

There are multiple sensors located within and outside the HVAC equipment to measure and monitor various environmental conditions and the operating status of the system and components. Some of the sensors can have user adjustable settings while others are non-adjustable and meant to keep the unit operating safely per the manufacturer. Examples of other sensors include smoke detectors, air filter pressure differential to indicate a clogged or dirty filter, exterior ambient air conditions, high or low power supply voltage and current, fan state, cooling status and stages, refrigerant temperature and pressure and many more. When a component fails, malfunctions or strays out of set operating boundaries (set points), alarms or alerts are displayed or automatically forwarded to a building monitoring system. During the following case studies, the data from devices and space conditions were obtained via direct IP connections or data aggregators and exported to a third-party monitoring system branded Foreseer® (Eaton). This system allows remote site visibility and device alarming, so we can track and trend space conditions and device operational status.

The example power dashboard shown in Figure 4 shows the various areas of power monitoring available at this site including service entrance, generators, direct current (DC) plant, fuse panels, HVAC with environmental relative humidity (RH), and 3<sup>rd</sup> party aggregators:



**Figure 4 - Power Monitoring Dashboard Example From Foreseer®**

In Figure 5 below we see the actual Return and Supply air temperatures as displayed at the CRAC unit and informed by on-board sensors along with room and outside air (OA) temperature. One can see that the cooling and room setpoint of 71°F is activating the cooling mode, as the recorded temperature in the room is 73.3°F -- indicating that the likely dead band setting is 2°F. Note that the Delta T (or  $\Delta T$  is the difference in Temperature) between supply and return temperatures is only 5.7°F and the mechanical cooling is operating only Stage 01. Typical discussion of set points at this site would be described as: “temperature set points are 71°F/2°F and RH set points are 40%/10%,” indicating set point and deadband setting.

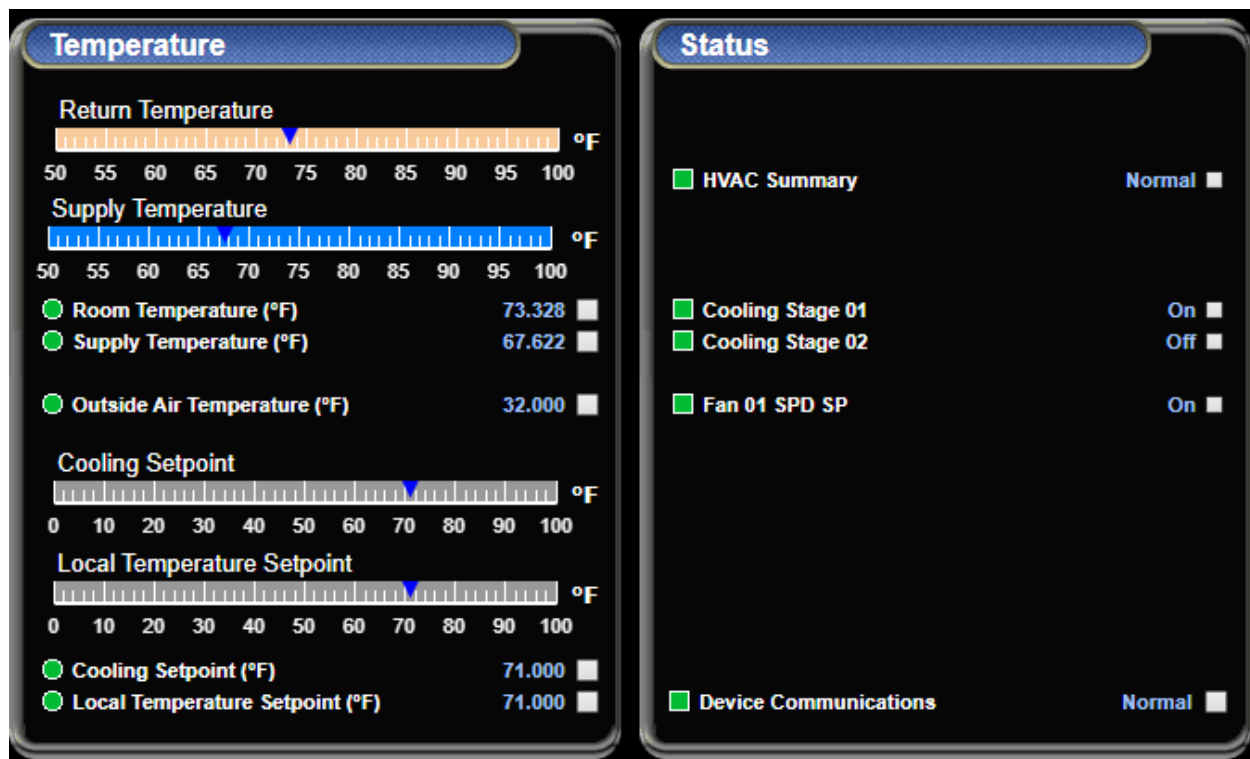



Figure 5 - HVAC Set Point Monitoring And Operational Display From Foreseer

### 3.1.2. ASHRAE Standards And Operating Ranges

The cable and data center industries, as well as state and local municipalities, have adopted many code requirements from various governing or advisory bodies. Of specific reference to this case study is the ASHRAE (American Society of Heating Refrigeration and Air-conditioning Engineers) 90.1-2019 “Energy Standards for Buildings Except Low-Rise Residential Buildings” and subsequent 90.4 -2019 “Energy Standard for Data Centers.” Generally accepted addendums also include publications from a formal ASHRAE Technical Committee (TC) TC 9.9, comprised of ASHRAE members who are subject matter experts in HVAC and data centers. Although headends are not yet called data centers, by definition many meet the criteria: Having a conditioned space, greater than 20Watts per square foot, and a production equipment load (commonly referred to as IT load) of more than 10 kW (ASHRAE, 2019). This widely accepted standard provides tables with ratings for various electronics equipment and the recommended safe operating temperature and humidity ranges (A1-A4), as shown in Figure 6.

In this case study, it was identified that many sites were operating at cooler than necessary space temperatures. Per the ASHRAE guidelines, the *recommended* area of the standard indicates that the temperature range is 18°C (64°F) to 27°C (80.6°F). Let it be noted that the *allowable* range is 15°C (59°F) at 80% RH up to 32°C (89.6°F) at 40% RH. This allowable range indicates the conditions that the production equipment can operate at occasionally, but only for limited transgressions and duration. These temperature and humidity ranges are meant to be measured at the face of the rack, where airflow is drawn in by the IT/Production equipment. The recommended ranges are the basis for this pilot, given that the premise is that raising the setpoints closer to the upper end of this range can be safely achieved, while significantly reducing energy costs, with little or no implementation or maintenance costs.



**Figure 6 - 2015 Recommended And Allowable Envelopes For Air-Cooled Equipment**

©ASHRAE [www.ashrae.org](http://www.ashrae.org), *Equipment Thermal guidelines for Data Processing Environments, Fourth Edition* (2015)

### 3.1.3. Capacity Impact

All major HVAC manufacturers test and rate their devices at various operating conditions, and both openly publish that data and include it in design specifications. Per their testing, the higher the operational limit is to a device, the greater the capacity of the unit. As shown in Figure 7, data from the Vertiv™ system design catalog has the lowest operational limit published, of 75°F, or 27.5 kW. If the space is maintained at this lower limit, then the operator is reducing the sensible kW capacity of the unit to approximately 20% lower than the same output at 85° (34 kW). If we follow the shown reductions of almost 2 kW per 5°F then a set point of 70°F would have a capacity of approximately 26 kW, or a 24% difference from 34 kW. This is stranded capacity and forces a de-rating of the units when capacity models are run. Further, it will require additional HVAC units brought online or purchased in order to match the heat rejection capability of the production load. This de-rating does not include the many other factors that could also reduce the capacity of the HVAC unit, such as poor airflow distribution, low Delta T and low humidity levels, all of which are great discussions for another paper.

## 2.1 Air-cooled Units Application and Performance Data

Table 2.1 Application data—Air-cooled, under-floor discharge<sup>1</sup> with EC fans

| Model Size                                                                           | 028             | 035             | 042             | 053             | 070             | 077             | 105              |
|--------------------------------------------------------------------------------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|------------------|
| <b>Net Application Data kBTUH (kW), Standard Air Volume and Evaporator Fan Motor</b> |                 |                 |                 |                 |                 |                 |                  |
| <b>Semi Hermetic Compressors (4 Step Cooling) with EC Fans</b>                       |                 |                 |                 |                 |                 |                 |                  |
| 85°F DB, 64.5°F WB, 52.3°F DP (29.4°C DB, 18.1°C WB) 32.4% RH                        |                 |                 |                 |                 |                 |                 |                  |
| Total kBTUH (kW)                                                                     | 116.4<br>(34.1) | 143.6<br>(42.1) | 163.6<br>(47.9) | 210.6<br>(61.7) | 247.1<br>(72.4) | 270.3<br>(79.2) | 352.5<br>(103.3) |
| Sensible kBTUH (kW)                                                                  | 116<br>(34)     | 141.2<br>(41.4) | 160.9<br>(47.9) | 210.2 (61.6)    | 245.9<br>(72)   | 269.2<br>(78.9) | 337.5<br>(98.9)  |
| 80°F DB, 62.9°F WB, 52.3°F DP (26.7°C DB, 17.1°C WB) 38.2% RH                        |                 |                 |                 |                 |                 |                 |                  |
| Total kBTUH (kW)                                                                     | 110.5 (32.4)    | 136.8 (40.1)    | 156.1 (45.7)    | 198.9 (58.3)    | 234.8 (68.8)    | 257.1 (75.3)    | 338.2 (99.1)     |
| Sensible kBTUH (kW)                                                                  | 105.9 (31)      | 128.1 (37.5)    | 146.2 (45.7)    | 193.9 (56.8)    | 226.6 (66.4)    | 248.8 (72.9)    | 307.2 (90)       |
| 75°F DB, 61.1°F WB, 52.3°F DP (23.9°C DB, 16.2°C WB) 45.1% RH                        |                 |                 |                 |                 |                 |                 |                  |
| Total kBTUH (kW)                                                                     | 105.3 (30.9)    | 130.8 (38.3)    | 149.5 (43.8)    | 189 (55.4)      | 223.9 (65.6)    | 245.8 (72)      | 325 (95.2)       |
| Sensible kBTUH (kW)                                                                  | 93.9 (27.5)     | 113.8 (33.3)    | 129.9 (43.8)    | 173.3 (50.8)    | 201.3 (59)      | 221.5 (64.9)    | 80 (273.1)       |

Figure 7 - Vertiv Application And Performance Comparison

### 3.1.4. Set Point Modifications And Results

Figure 8 is a screenshot of a Vertiv CRAC unit with a return temperature set point of 67°F and sensor readings of 68°F and 39% RH. The fan is operating at 81% and mechanical cooling is at 72% of capacity.





Figure 8 - Vertiv Performance And Set Point Display With Function Key

Nine sites were selected for this pilot (see Table 1) in conjunction with adding lockable and programmable thermostats to control the room temperatures. Utility incentives were awarded by Com Ed



to supply and install the thermostats and raise the set points. Energy use data was recorded and averaged 30 days prior to any changes (August 2019) and averaged for 30 days after the adjustments (September 2019). All sites were within 40 miles of each other, to maintain consistency of local weather impacts. The typical CRAC settings we encountered on this project had temperature set points of 68°F with a 2°F deadband. Humidity was 40% RH with a 10% deadband. The set points were raised by approximately 1°F per day over 2 weeks. There were no recorded thermal events or issues during this time. The average outdoor temperature during August 2019 = 74°F and RH = 68%  
The average outdoor temperature during September 2019 = 71°F and RH = 75%

**Table 1 - Site Energy Reduction Percentages After Set Point Increase**



### Summary

- Pre-adjustment energy savings estimate for this pilot was 5% of HVAC energy per site. Actual results averaged 20% of HVAC kWh savings.
- Addition of lockable thermostats allowed for a 60% project cost reduction via utility incentive
- Actual cost to adjust set points = \$0 (self-performed by Comcast engineers)
- To date thermal events or impact at these sites with raised set points = 0
- Results of this pilot exceeded expected results

### 3.2. Blanking Panel Installation

Blanking panels (BPs) are a generic term for any device or product that will block open rack unit (RU) spaces in a typical production load cabinet, to prevent cold air from by-passing the in-rack devices. This by-pass conditioned cold air is considered wasted cooling, as it did not enter any production equipment and will return to the CRAC units unused. The BP is meant to plug the open rack gap and reduce the wasted cooling resource. See Figure 9 below. Blanking panels are an essential part of airflow management in a headend or data center and considered low hanging fruit to save energy and improve airflow distribution.



**Figure 9 - Open RUs On Left, PlenaForm Systems Blanking Panels Installed On Right**

### ***3.2.1. Types Of Blanking Panels***

There are various sizes of pre-manufactured products fitting 1,2,4,6 and 10 RU in a single piece of material. The material is generally a fire rated molded plastic, sheet stock or sheet metal panels. One of the most common rack sizes is 42 RU, any time there is an open RU space, a blanking panel should be installed to maintain industry best practices for airflow efficiency. Common product types are shown below in Figures 10 and 11.



**Figure 10 - Upsite Technologies 2 RU White Molded Blanking Panel**

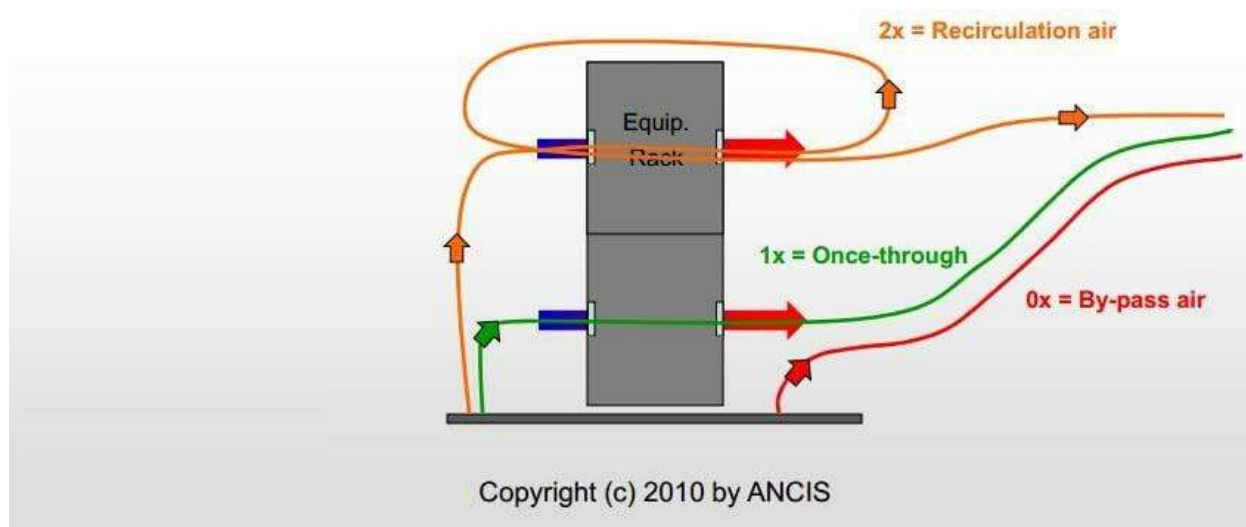


**Figure 11 - Polargy PolarFlex Sheet Blanking Panel**

### **3.2.2. Airflow Efficiency**

The visual depiction of airflow through and around an IT rack is well described below in Figure 12 which is still relevant and accurate to show where energy and efficiency is lost. This timeless summary by Dr. Magnus Herrlin (Herrlin, 2010) represents the good, the bad and the ugly from an airflow perspective. Cooling resources that are introduced into a space that move through an IT device (router, switch, modem or server) to perform work is the good. Airflow that never reaches its intended target (the in rack devices) is downright bad and that air that gets caught in an endless recirculating loop over the top of a rack, at an aisle end or through a gap where the BP should be is just plain ugly. Airflow management is comprised of several components and variables including rack layout, rack equipment orientation, CRAC positions, ceiling height and more, yet blanking panels are simply a must and diligence to this requirement is critical for efficiency. Many spaces today do not have full BP installations, yet the rooms stay mostly in compliance by brute force of the air conditioning (i.e. always running). It is not necessarily that the space needs more cooling tons but rather that the airflow distribution is so poor that many units are running with fan only. This has unintended consequences as the HVAC fan motor produces heat which is now added to the room. As well by operating in fan only mode the HVAC unit is pulling in warm return air and pushing this warm air back into the room without conditioning, meaning the supply air temperature is actually being raised by the HVAC unit.





*In the illustration above, from Herrlin's presentation titled Room-Level Energy and Thermal Management in Data Centers: The DOE Air Management Tool, you see that once-through cooling passes through the equipment rack only one time before returning to a cooling unit while recirculation air passes through twice. Bypass airflow on the other hand, does not pass through the equipment rack at all before returning to a cooling unit, hence bypass airflow.*

**Figure 12 - Visual Display of Airflow Anomalies**


### **3.2.3. Blanking Panel Case Study Results**

In this study we selected 7 sites across the US which were of varying sizes. We recorded the HVAC kW consumption before and after installation and examined consumption over the following 60 days. Weather changes were accounted for and data was averaged accordingly. One site in the west experienced an increase in energy use that was traced back to installation just before a prolonged and significant heat wave moved into this location. Labor to install was mostly self-performed with 3 sites using outside labor. Primary blanking panels used were the one and two RU molded plastic snap-in place devices.

#### **Summary**

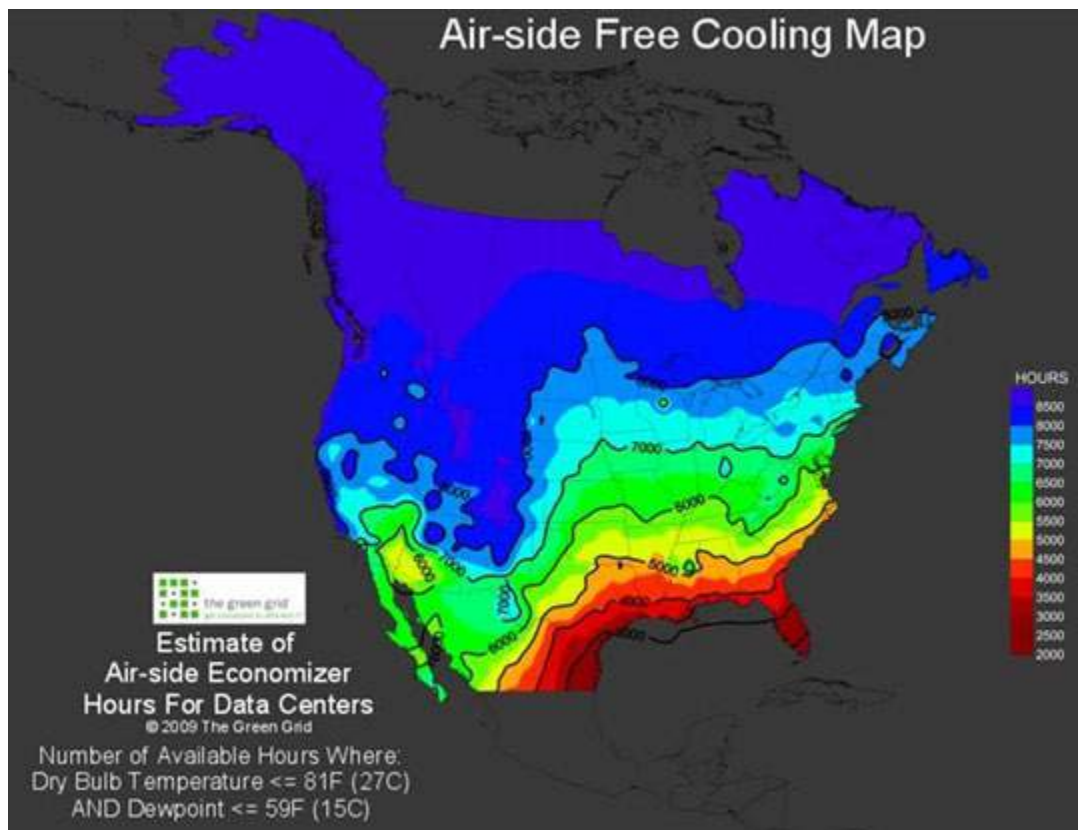
- Pre-installation projections for HVAC cooling energy savings after blanking panel installation was 3-5%.
- The average energy savings reported as calculated from data in Figure 13 below was 16%. If we remove both the high and low readings from the list our average savings is 12%.
- The installation of blanking panels does not require any special tools or knowledge.
- Low cost low tech with high impact to cost reductions and reliability

**Table 2 - Site Energy Reduction Percentages After Blanking Panels Installation**



### 3.3. HVAC Economizer Options

The term economizer has become a generalization for many of the technologies or components that make an HVAC device more energy efficient. These include variable speed drives (VSD) for fans and pumps, electronically commutated motors, dual circuits (refrigerant and glycol/water mix) at the condensers, variable compressors and direct outside air. The technology discussed here will pertain specifically to airside economization for packaged ground mount or roof top HVAC units (RTU), where ambient outside air is mixed or directly fed into the headend space when temperatures are favorable for use. Typically, this is an add-on option on many devices, but the investment is worthwhile as the payback period is quite short due to substantial energy savings achieved in most parts of the US and Canada as shown in Figure 14. With the exception of the states bordering the Gulf coast most of North America can support over 5000 hours of economizer operation which is 57% of the hours in a year (8,760 hrs annually). This is a significant reduction of time required to run mechanical cooling which reduces operational expenses (OpEx) and extends the life of equipment.



**Figure 13 - Hours With Ideal Conditions For Air-side Economization (DOE And The Green Grid)**

### 3.1.1 Economizer Operations


Most of the HVAC manufacturers combine the control operations of the Direct Expansion (DX) packaged HVAC device with the economizer components to provide for an automatic engagement when the outside weather is favorable or meets the requirements of the economizer set points (See chapter 1.4 of this paper). The economizer control setpoints will identify what the ambient exterior conditions need be to activate the dampers and louvers for fresh air mixing. Direct introduction of cold outside air into a space can create problems so the economizer should be mixing warm return air from the space with the colder air so thermal shock, low temperature alarms or condensation does not occur. This blending is necessary during colder winter months. Many geographical locations experience long term ideal conditions for economization and can position the set points for temperature, humidity, enthalpy or dewpoint control. It is typically the interior temperature set point that is the main controller or restrictor to full economization capability. If the set points for temperature and humidity are set higher than the economizer hours can be extended. Many manufacturers are using dew point temperatures as a more accurate and reliable control point for higher efficiency of their HVAC systems. The industry now recognizes that by allowing a wider range of RH, and given proper controls, a great deal of energy and water can be saved while maintaining acceptable IT performance. (Sorell, 2017). Simply stated the control system activates dampers to allow outside air into the device supply duct assembly providing cooling to the space and exhausts the hot room air directly outside. This can occur on one or all of the HVAC units per the sequence of operation set by the end user. As shown in figure 14 below of a sample HVAC unit dashboard the outside air (OA) is sufficiently cold for economization and the relative humidity is also sufficient to allow for economization.

The unit is allowing the OA damper to open to 46% mixing it with warmer return air to bring the supply temperature up to 54.3°F. The set point (return air temperature) for this space is set at 68.2°F which is unnecessarily low thus reducing many potential hours of economization.



**Figure 14 - Aaon Dashboard With Operational Performance Of Economizer**


Figure 15 below shows a comparison of average Amps without Economizer engaged 55.6A (red) and with economizer engaged 28.9 (green) which results in a 48% energy reduction. The orange bottom line indicates when the economizer cycled on and off in October and November. This location seasonal climate change provides for 5 to 6 months of economization hours at a minimum when the control set point for activation of the economizer is set at 50° F.



**Figure 15 - Energy Use Comparison With And Without Economizer Activated**

As shown in Table 3 the average energy savings projected across 10 sites is 38%. This figure can be increased by adjusting ambient room temperature control set points and the economizer set points. The MA 2 site indicated a reduction of 58%. These values seemed quite aggressive and with handheld meters and instruments our electrician confirmed the amp readings as correct.

**Table 3 - 5 Month Projected Seasonal Energy Savings From Economizer Engagement**



**Summary:**

- Most of North America is suitable for economization use
- Automated system controls require no manual intervention
- Significant energy reduction is achieved when economization is engaged
- Some state codes require economizer feature on HVAC

## 4. Conclusion

These three case studies have allowed us to identify substantial measured savings from multiple energy conservation measures. Although considered best practices blanking panels and many airflow efficiency measures often get overlooked or given low priority until an alarm indicates a thermal event. In many instances the network technicians who are “rack and stacking” their equipment in the space are not fully aware of the cost and impact of good airflow management resulting in limited compliance. On the other hand, those same technicians are keenly aware when a space is un-necessarily cold. It is not uncommon to see network or cabling tech’s wearing coats inside a headend with low set points and we can surmise that not only is that space wasting energy but the productivity of the workers are likely to be reduced as well. The economizer technology available today has become an outstanding energy saver and has not received the attention or accolades that should make it a requirement everywhere. Several states have adopted building codes and included HVAC economizers as mandatory, but the magnitude of savings should support inclusion regardless of codes. It is the intention of this paper to raise awareness to the potential for significant verifiable energy cost reductions and associated emissions that if deployed and measured across the cable industry would help achieve Energy 2020 goals quite easily. Susan JinDavis VP of Environmental Affairs and Chief Sustainability Officer of Comcast recently stated “Different energy strategies are needed for our various locations. We are choosing solutions based on what will make the biggest difference location by location, whether that’s on-site solar, green tariffs, renewable energy

supply contracts, or a combination. Across all of our operations, we want to find ways to not only power with renewables, *but to reduce our power needs overall*”. As stated in the introduction the industry as a whole has been self-directed to maintain a commitment for efficiency and sustainability and the results of the measures presented here can provide financial justification for inclusion in long term planning and practices.

## Abbreviations

|         |                                                                          |
|---------|--------------------------------------------------------------------------|
| AFO     | Airflow Optimization                                                     |
| ASHRAE  | American Society of Heating Refrigeration and Air-conditioning Engineers |
| CFM     | Cubic feet per minute (airflow measurement)                              |
| Com Ed  | Commonwealth Edison (Utility company)                                    |
| CRAC    | Computer room air conditioner                                            |
| CI      | Critical Infrastructure                                                  |
| Delta T | Difference in temperature                                                |
| DF      | Down Flow                                                                |
| DP      | Dew point                                                                |
| ECM     | Energy Conservation Measure                                              |
| F       | Fahrenheit                                                               |
| HVAC    | Heating Ventilation and Air Conditioning                                 |
| ISBE    | International Society of Broadband Experts                               |
| OpEx    | Operational Expense                                                      |
| RA      | Return Air                                                               |
| RTU     | Roof Top Unit                                                            |
| RH      | Relative Humidity                                                        |
| SA      | Supply Air                                                               |
| SCTE    | Society of Cable Telecommunications Engineers                            |
| SNE     | Small Network Equipment                                                  |
| STB     | Set Top Boxes                                                            |
| TC      | Technical Committee                                                      |
| VFD     | Variable Speed Drive                                                     |
| VSD     | Variable Frequency Drive                                                 |

## Bibliography & References

- 1 Page 7- United States Data Center Usage Report No. LBNL-1005775, Lawrence Berkley National Lab-Energy Technologies Area, Arman Shehabi et al. 2016
- 2 Page 8- Energy Conservation Measure Recommendations for Cable Edge Facilities, A Technical Paper prepared for SCTE/ISBE by Daniel Marut et al.2017
- 3 Page 10- ASHRAE Standard 90.4- 2019, Energy Standard for Data Centers (page 4 definitions), Richard Zbin (Chairman) et al. 2019.

- 4 Page 15- Room-Level Energy and Thermal Management in Data Centers: The DOE Air Management Tool. Presentation at IMAPS Advanced Technology Workshop on Thermal Management Palo Alto, CA, September 30, 2019. Magnus K. Herrlin Ph.D.
- 5 Page 18- Humidity Control in Data Centers, Center of Expertise for Energy Efficiency in Data Centers, Lawrence Berkley National Lab, Vali Sorell and Magnus Herrlin, 2017



# Machine Learning Techniques for Equalizing Nonlinear Distortion

A Technical Paper prepared for SCTE•ISBE by

**Rob Thompson**

Director, NGAN Network Architecture  
Comcast  
1800 Arch St., Philadelphia, PA 19103  
(215) 286-7378  
robert\_thompson6@cable.comcast.com

**Xiaohua Li**

Associate Professor, Dept. of Electrical and Computer Engineering  
State University of New York at Binghamton  
Binghamton, NY 13902  
(607) 777-6048  
xli@binghamton.edu

# Table of Contents

| Title                                                                   | Page Number |
|-------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                    | 4           |
| 2. Artificial Intelligence (AI).....                                    | 4           |
| 2.1. Historical Perspective .....                                       | 4           |
| 2.2. Common Solutions .....                                             | 4           |
| 2.3. Popular Tools/Models .....                                         | 5           |
| 2.4. Biological Inspiration.....                                        | 5           |
| 2.5. DOCSIS Transmit Pre-Equalization .....                             | 6           |
| 3. Power Amplifier (PA) Efficiency Problem .....                        | 6           |
| 4. Nonlinear Distortion (NLD) .....                                     | 7           |
| 4.1. Digital Signal Impact.....                                         | 8           |
| 4.1. Adjacent Channel Leakage Ratio (ACLR) Measurements .....           | 11          |
| 4.2. PA Industry.....                                                   | 13          |
| 4.1. NLD Modeling.....                                                  | 14          |
| 5. NLD Mitigation .....                                                 | 18          |
| 5.1. Peak-to-Average-Power-Ratio (PAPR) Reduction .....                 | 18          |
| 5.1. Transmitter Digital Pre-Distortion (DPD).....                      | 23          |
| 5.1. Receiver Post-Distorter Equalization.....                          | 28          |
| 6. Severe NLD Mitigation .....                                          | 28          |
| 6.1. Enhanced Equalization – Integrating Volterra Series and DNNs ..... | 29          |
| 7. Conclusion .....                                                     | 34          |
| Abbreviations.....                                                      | 34          |
| Bibliography & References .....                                         | 36          |

## List of Figures

| Title                                                                                  | Page Number |
|----------------------------------------------------------------------------------------|-------------|
| Figure 1 - Biologically Inspired Perceptron Model.....                                 | 5           |
| Figure 2 - DOCSIS Transmit Pre-Equalizer Structure .....                               | 6           |
| Figure 3 - Power Output and Gain Compression Characteristics of a PA.....              | 7           |
| Figure 4 - Sample RPD Node .....                                                       | 8           |
| Figure 5 – Error Vector Magnitude (EVM) for a QPSK Signal .....                        | 10          |
| Figure 6 - Output Spectrum Comprising Allocated and Adjacent Channels.....             | 11          |
| Figure 7 - Increasing Spectral Regrowth over Amplifier Linear Region .....             | 12          |
| Figure 8 - 3 Tone NLD Measured vs. Simulated.....                                      | 18          |
| Figure 9 - PA Characteristics with PAPR and DPD.....                                   | 19          |
| Figure 10 - DOCSIS Downstream SC-QAM PAPR Measurement.....                             | 21          |
| Figure 11 - Traditional Selected Mapping, PAPR Reduction via Pre-Coding .....          | 22          |
| Figure 12 - Transformation Based Pre-Coding for SLM PAPR Reduction .....               | 22          |
| Figure 13 - High-Level Transmitter with CFR and Adaptive DPD.....                      | 23          |
| Figure 14 - Nonlinearized PA vs. Linearized PA via Digital Pre-Distortion.....         | 24          |
| Figure 15 - CM Upstream Transmit Pre-Equalization and Post-Equalization Functions..... | 25          |
| Figure 16 - Effects of Sample Rate on a Pre-Distorted Signal.....                      | 26          |
| Figure 17 - Memory and Memoryless DPD Results .....                                    | 27          |

|                                                                                                         |    |
|---------------------------------------------------------------------------------------------------------|----|
| Figure 18 - Block Diagram of DNN Equalizer .....                                                        | 29 |
| Figure 19 - System Block Diagram with Nonlinear Power Amplifier and Deep Neural Network Equalizer ..... | 30 |
| Figure 20 - Constellation of 16-QAM Non-Equalized vs. Equalized.....                                    | 32 |
| Figure 21 - Comparing Three Equalization Methods for 16-QAM under Various NLD Levels .....              | 33 |

## List of Tables

| <b>Title</b>                                                                       | <b>Page Number</b> |
|------------------------------------------------------------------------------------|--------------------|
| Table 1 - Sample RPD Node Datasheet Summary .....                                  | 9                  |
| Table 2 - PA Classes .....                                                         | 13                 |
| Table 3 - Comparing MSE/SER Improvement % for the Three Equalization Methods ..... | 33                 |

# 1. Introduction

Since the early 2000s, the cable television (CATV) industry has been playing its part in the Artificial Intelligence (AI) community by deploying equalization technology to enable its digital signals to survive varying frequency response conditions within its cable plants. Simon Haykin describes how the perceptron and the adaptive filter using the least mean squares (LMS) algorithm are naturally related [1]. Equalization has evolved into a powerful tool, enabling the CATV industry to achieve communication efficiencies once thought impossible -- but that story is not quite complete. The limits of equalization may extend beyond the linear frequency response, and cancel the nonlinear responses commonly associated with nodes and other active devices which use power amplifiers (PAs). Achieving nonlinear equalization requires new equalization methods, like receiver post-distorter equalization, where techniques include AI models, such as deep neural networks (DNNs). Furthermore, researchers have been advancing nonlinear distortion cancellation via other methods, including peak-to-average-power-ratio (PAPR) reduction, and digital pre-distortion (DPD). These technologies are beginning to show up in newer generation devices, where demands for radio frequency (RF) output power is high, while keeping power consumption low, like the full duplex DOCSIS (FDX) remote PHY device (RPD) nodes. DPD technologies cancel the contribution of the transmitting device only. More aggressive nonlinear distortion cancellation methods may be accomplished by advanced DNN approaches, such as incorporating input features derived from Volterra series models, which has become a popular model for nonlinear distortion that can be used to describe multiple nonlinearity orders and memory. Then efficiencies across the CATV network could be considered, either by higher node RF output power, or more efficient PA architecture/bias within the node, amplifier, and/or customer premise equipment (CPE). This paper will propose how current CATV equalization systems could be enhanced to cancel severe nonlinear distortion based on some of these novel approaches to nonlinear equalization.

## 2. Artificial Intelligence (AI)

### 2.1. Historical Perspective

AI has existed for a very long time -- close to 80 years. In 1943, Warren Sturgis McCulloch and Walter Pitts published a paper titled “A Logical Calculus of Ideas Immanent in Nervous Activity,” laying the foundations for artificial neural networks (ANNs) [1]. Since then, many ideas involving AI have been shared, and this community has grown appreciably. Patrick Winston, who was born in 1943 and later became a MIT professor who taught a course in AI, described it as being about algorithms, enabled by constraints, exposed by suite of representations, that support the development of models targeted at thinking, perception, and action [2]. That definition is inclusive of many things -- in fact, some very simple internet searches can yield timelines rich with AI milestones, including events such as when Deep Blue defeated World Chess champion Garry Kasparov in 1997, or a more recent milestone, on October 15th, 2019, when OpenAI enabled a robot to learn how to single-handedly solve Rubik’s Cube with the support of two neural networks [3].

### 2.2. Common Solutions

There are many artificially intelligent solutions that we encounter every day, possibly without even realizing it. Comcast, for example, provides multiple products which incorporate AI technology. Some of these solutions include the Voice Remote, that adapts to the uniqueness of how each and every one of us speaks, and in turn assists with accessing and enjoying content on Comcast’s X1 platform. The X1 recommendation engine detects patterns in the content we consume, and assists users in navigating a wealth of available content and offer recommendations. Internet-based products like xFi Advanced


Security protect our home networks against constantly evolving network threats. Interactive assistants, like xFinity Assistant, help customers by leveraging an interactive knowledge base of common solutions.

### 2.3. Popular Tools/Models

There are many available AI tools and frameworks, including TensorFlow, and PyTorch [4]. These systems are designed to assist with navigating the vast array of models, each with their unique set of pros and cons, when it comes to approximating the functions that couple input and output data patterns together. Some of these tools readers may have already heard of, like DNNs [5]. Others, like support vector machines (SVMs), may be less familiar. Fortunately, finding the right model fit for a particular problem has been automated via tools like automated machine learning (Auto ML), which not only selects the best function approximation model, but also assists with tuning the parameters of that model to optimize its training and generalization properties.

### 2.4. Biological Inspiration


The perceptron model gets its inspiration from the Pyramidal cell shown Figure 1 [1]. One of the key characteristics of this model are its synaptic weights, which it applies to each of its input signals. The input signals are summed together and applied to an activation function. The sigmoid function is a popular activation function, which limits the output response to a specific range of continuous values. An Artificial Neural Network (ANN) connects multiple perceptions together in a variety of ways, in parallel, and/or in series. In doing so, interesting behaviors begin to emerge -- the most interesting being the nonlinear adaptation of the model weights. This could present interesting opportunities for linear adaptive equalization systems used today, where nonlinear adaptation enhancements could enable systems to account for both linear and nonlinear distortion present within the communication channel to improve the performance of equalization systems overall.



**Figure 1 - Biologically Inspired Perceptron Model**

## 2.5. DOCSIS Transmit Pre-Equalization

Since the early days of the Data Over Cable Service Interface Specifications (DOCSIS), equalization has enabled the cable modem termination systems (CMTSs) to adapt and learn the unique frequency response shared between it and each of the cable modems (CMs) to which it is connected. The CMTS shares this knowledge with the CM, via a coefficient vector or weights often times referred to as “taps”, asking it to either convolve or overwrite its current set of weights, based on how quickly the frequency response was changing. The CM applies these weights to future transmissions to the CMTS, to cancel the frequency response effects of the channel, which could be micro reflections (echoes) or filter effects including group delay or amplitude roll-off [6]. This form of equalization came to be known as “transmit pre-equalization” in DOCSIS 1.1 [7].



**Figure 2 - DOCSIS Transmit Pre-Equalizer Structure**

Comparing the DOCSIS 2.0 equalizer of Figure 2 to the perceptron of Figure 1, one cannot help but notice the similarities between the two models. What is most like the perceptron is the linear adaptive filter’s weighted inputs feeding a linear combiner, and the ability to perform continuous learning – a single neuron operating in its linear mode [1]. The perceptron and an adaptive filter using LMS are naturally related [1].

## 3. Power Amplifier (PA) Efficiency Problem

Designer Charles Warren once gave a talk titled “How Might We: Three Words That Make Design Better” [8]. His thoughts are both refreshingly entertaining, and very helpful in organizing our ideas around innovation and establishing the following goal statement for this paper.

*How might we optimize PA efficiency in our RPDs?*


Improving PA efficiency in our RPDs may be beneficial in maintaining existing requirements, like RF output levels, as we introduce new capacity-enhancing technology. FDX falls into this category. Its accompanying echo cancellation (EC) technology, is necessary for facilitating bidirectional communication at the same operating frequencies [7]. EC technology additions will require compromises, especially while maintaining existing requirements for RPD RF output power, complexity, weight, power consumption, heat dissipation and cost.

At this point, the reader may be thinking “The goal statement is limited to RPDs, but why not optimize all the active components within the RF chain, including line extenders, mini-bridgers, trunk amplifiers, home drop amplifiers and even CPE front ends?” This is a thought we hope to address in paper as well, but in the spirit of following Warren’s “How Might We” process through to its end, let’s first consider some of the things that may stopping us from achieving our goal as stated.

First and foremost, there is a catch when it comes to increasing PA efficiency, and that is increased nonlinearity. Any savings in PA power consumption will have to be balanced with a system of mitigating any increases in nonlinearity, something we intend to explore fully in this paper. Another consideration is with respect to standards, and whether the solution requires standardization to ensure seamless interoperation across vendors that can be deployed, operated, and maintained in a consistent manner. Lastly, with introduction of any new technology, there is always consideration given to how to gracefully coexist with legacy products and services, ideally minimizing any impact to existing services.

## 4. Nonlinear Distortion (NLD)

Figure 3 illustrates how PAs strengthen their input signals [10]. When a PA's input power is at its lowest levels, its output power behavior is more linear than it is nonlinear, and its gain is constant. Ideally, PAs would behave linearly for all input signal levels, including high input powers, as illustrated by the dotted line. However, practical PAs generally available today cannot strengthen input signals without adding nonlinear distortion (NLD) to those input signals. As we will later see, NLD increases more rapidly than the illustrated input/output increases of the fundamental signals. Eventually, the PA reaches saturation, and its performance becomes more nonlinear than linear. At this point, the PA's output is no longer proportional to its input, and its performance is dominated by NLD. Further, a PA's linear operating region or dynamic range is a range of input powers that include a predictable mixture of impairments, including noise and NLD. At low input power, noise dominates the impairment mixture, but as input power increases, noise performance improves, while NLD worsens. The challenge for the network designer is to strike a balance between noise and NLD, so that their combined performance is within acceptable limits. Output-power-back-off (OBO) is a term used to describe this compromise, where the PA's operating point is typically several decibels (dBs) below its compression point and includes acceptable noise and NLD levels for overall system performance [11].



**Figure 3 - Power Output and Gain Compression Characteristics of a PA**

## 4.1. Digital Signal Impact

Metrics pertaining to Figure 4 node downstream transmission path (blue) performance for noise and NLD have been included in Table 1 [9]. These specifications are measured using 6 MHz wide channels and are based on a full channel loading consisting of 194 single carrier quadrature amplitude modulation (SC-QAM) signals. Decibel-millivolts, dBmVs, are generally used, for mathematical convenience, within the CATV industry, to reference operating levels,  $O$ , for analysis of these systems [12].

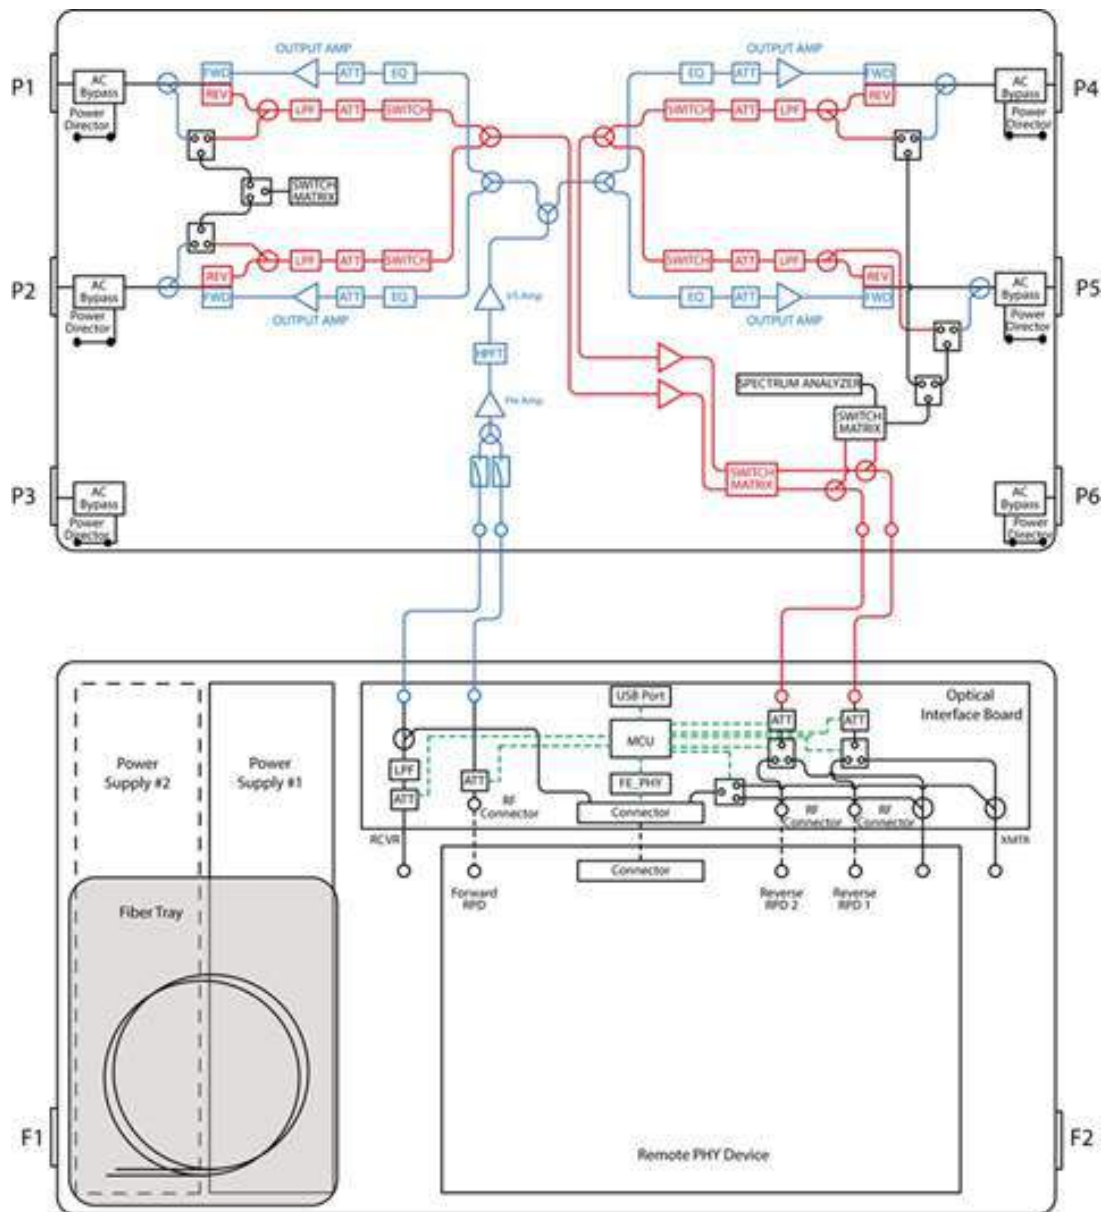


Figure 4 - Sample RPD Node



**Table 1 - Sample RPD Node Datasheet Summary**

| Sample RPD Node Specifications               |                               |
|----------------------------------------------|-------------------------------|
| Minimum operational gain, $G$                | 42 dB                         |
| Noise Figure, $NF$                           | 15.5 dB at 54 MHz             |
| Composite-Intermodulation-Noise, $CIN_{ref}$ | 50 dB                         |
| Reference output level, $O_{ref}$            | 42 dBmV at 54 MHz             |
| Power                                        | 160.6 W @ 2.16 A, and 90 V AC |
| Weight                                       | 49.8 lbs.                     |

Signal-to-noise ratio (SNR) describes the relative measure of signal power,  $S$ , to the noise power,  $N_p$ , which is the thermal noise or noise floor measured within the same bandwidth as the signal  $S$ , in this case 6 MHz.  $N_p$  is estimated using (1), where  $k$  is Boltzman's constant ( $1.374 \times 10^{-23}$  joules/ $^{\circ}K$ ),  $T$  is the absolute temperature in degrees Kelvin ( $^{\circ}K$ ), and  $B$  is the bandwidth of the measurement in Hertz (Hz).  $N_p$  in the CATV industry is typically expressed in terms relative to 1 milli-volt (mV) across a  $75 \Omega$  impedance, therefore  $N_p$  at  $62^{\circ}F$  is approximately -57.4 dBmV [13].

$$N_p = kTB \quad (1)$$

Composite Intermodulation Noise (CIN) is a type of NLD which results from nonlinear distortion generated from loading conditions, which include digital signals, like SC-QAM. Node contribution for SNR and CIN can be calculated using (2) and (3) respectively, for changes in its output levels  $O$  constrained over the node's dynamic range [12], [13]. In the CATV industry, CIN is typically dominated by 3rd-order NLD, which may not be the case for other communication systems, like those used in the satellite industry [14].

$$SNR = O + 57.4 - G - NF \quad (2)$$


$$CIN = CIN_{ref} - 2(O - O_{ref}) \quad (3)$$

Increasing the node's operating output power, from the originally specified  $O_{ref}$ , by 5 dB, will result in a new output level,  $O = 47$  dBmV at 54 MHz. This new output level increase will also increase SNR to 47 dB at 54 MHz, using (2). However, CIN in (3) will decrease to 40 dB, leading to a 2 dB degradation overall of the System SNR,  $SNR_S$ , per (4).

$$SNR_S = -10 \log_{10} \left( 10^{-\frac{SNR}{10}} + 10^{-\frac{CIN}{10}} \right) \quad (4)$$

Therefore, increasing node RF output signal levels may enable the designer to improve homes-per-node efficiency, such as with higher output level  $O$ , but will do so at the expense of increasing the node's power consumption and degrading the overall system performance criteria,  $SNR_S$ . Increasing PA RF output levels of networked devices is one of the ways in which to optimize PA efficiency. However,

higher RF output levels may drive power consumption above the network operator's acceptable threshold -- in the case of the node example, above a 160 W maximum. Power consumption threshold values are based on network operator's unique powering constraints, which may be limited by multiple factors, including network design, hardware capability and local regulatory restrictions. Regulatory restrictions here specifically involve the placement of powering hardware at specific telephone pole or pedestal locations. Increased power may also impact heat dissipation and design of the node's housing. These impacts may translate to the larger surface area and weight of the node's housing to facilitate necessary heat transfer, where weight may increase above an acceptable threshold, in this case 50 lbs. maximum. Node weight is an important factor, because technicians need to be able to lift the node in order to connect it to the cable plant. The activity could be above ground, attached to telephone poles, or below ground within a pedestal mount.



**Figure 5 – Error Vector Magnitude (EVM) for a QPSK Signal**


The effect of degraded performance on digital signals is illustrated in Figure 5. A quadrature phase shift keying (QPSK) signal degraded by  $SNR_S$ , will cause actual symbol reception to deviate from the ideal symbol receive point, shown as a dark circle. The resultant error vector,  $e_j$ , between the actual symbol and the ideal QPSK symbol receive point represents a measure of fidelity. Modulation error ratio (MER) in (5) measures the cluster variance in dB, that can be observed in a SC-QAM signal. It includes the effects of inter-symbol interference (ISI) spurious, phase noise, and all other degradations, where  $E_{av}$  is the average constellation energy for equally likely symbols, and  $N$  is the number of symbols averaged [7].

$$MER_{symb}(dB) = 10 \log_{10} \left\{ \frac{E_{av}}{\frac{1}{N} \sum_{j=1}^N |e_j|^2} \right\} \quad 5$$

Poor MER can lead to symbol decision boundary crossings, translating to symbol errors. If frequent enough, these symbol errors can overwhelm forward error correction (FEC) schema, leading to packet errors and ultimately loss of network payload.

#### 4.1. Adjacent Channel Leakage Ratio (ACLR) Measurements

Figure 6 illustrates a nonlinear response to a band-limited signal, where the output of a system represents both linear (solid line) and nonlinear components (dashed line) [10]. The term intermodulation distortion (IMD) is synonymous with CIN.




**Figure 6 - Output Spectrum Comprising Allocated and Adjacent Channels**

Figure 6 is an example of how CIN can accumulate under the signal(s). One in-band measurement approach of CIN involves the following steps:

1. Measuring  $SNR_S$ , in dBmV per 6 MHz
2. Turning the signal load off
3. Measuring the thermal noise contribution to  $SNR_S$  at the same frequency,  $N_P$ , also in dBmV per 6 MHz
4. Calculating the difference between these two values, via algebraic manipulation of equation 4, as the CIN contribution

This is a reasonable measurement approach in a lab environment, where turning the PA's signal load off will likely not negatively impact a customer. Figure 6 illustrates a more customer-friendly approach that involves out-of-band distortion measurements of the integrated power in the adjacent channel(s), called adjacent channel power ratio (ACPR) or adjacent channel leakage ratio (ACLR), where these measurements are expressed in decibels relative to the signal's channel power, or -dBc [10].



**Figure 7 - Increasing Spectral Regrowth over Amplifier Linear Region**

Figure 7 is composed of a series of output measurements of one PA, where each color corresponds with 5 dB changes in ACLR or spectral regrowth, resulting from incremental adjustment of the PA input power. ACLR measurements were made via Keysight’s PXA signal analyzer model N9030A and vector signal analyzer (VSA) software, model 89601B. The 6.4 MHz bandwidth power delta marker measurements of both the received signal power and the spectral regrowth of the upper first adjacent channel were used to obtain the ACLR measurements.

Figure 7 data was obtained from a CATV drop amplifier, which is sometimes used within the customer’s home to overcome losses associated with providing signals to multiple (1-3) client devices, primarily supporting video services, but also including voice and high-speed data services. PAs are used in many CATV network clients’ front ends, used within the customer’s home, like set-top boxes (STBs), cable modems (CMs), and digital terminal adapters (DTAs). The outside plant represents another group of network elements with embedded PAs, including the node previously discussed and additional amplifiers, called trunk amplifiers, line extenders and bridging amplifiers, which are used to compensate for cabling and passive losses incurred while distributing and aggregating services to and from the customer’s home.

All of these cascaded PAs aggregate noise and NLD, resulting in an accumulated end-of-line (EOL) performance, which further restricts the individual contribution for any one PA in the chain to that of even higher fidelity (i.e. 50 dB per Table 1). This ensures that the customer’s services meet, or ideally exceed, some minimum service level agreement (SLA). Additionally, variations in RF levels can occur for multiple reasons, like temperature changes, wind loading, and plant maintenance. Sometimes RF levels can change appreciably over short periods of time. Network designers may specify even better performance, to accommodate this performance variation in several of the components or sub systems in the network chain, a.k.a. ‘margin-stacking’ [10].

Similar ecosystems exist in other industries, including cellular, Wi-Fi, satellite, and the Internet of Things (IoT). For example, massive MIMO and millimeter wave transmissions use many PAs in cellular deployments and will also have similar end-to-end performance requirements [11].

## 4.2. PA Industry

The PA industry has a growth outlook from 2019 to 2025, with compound annual growth rate (CAGR) of 7.6%, primarily from anticipated developments in newer generation cellular communications. Some of the key providers include NXP Semiconductor (Netherlands), Broadcom Corporation (U.S.), Qorvo Inc. (U.S.), Anadigics Inc. (U.S.), RFHIC Corporation (U.S.), TekTelic Communications Inc. (U.S.), Texas Instrument (U.S.) among others [27].

Providers can supply a diverse range of PA classes suited for a variety of applications, some of which have been summarized in Table 2Table 2 - PA Classes.

**Table 2 - PA Classes**

| Class     | Description                                                                                                                                                                                                                                                                                                                        |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| A         | Very linear, especially for smaller signal amplitudes. Used for PAs at high millimeter-wave frequencies, but considered too inefficient, < 50%, for PAs in cellular wireless communications applications                                                                                                                           |
| B         | Power transistor is biased just at the threshold voltage, so the transistor conducts for only a half-cycle of the RF waveform. This wave generates a significant amount of even harmonic distortion, and some self-bias. Maximum theoretical efficiency rises to 78.6%, compared with a Class A.                                   |
| AB        | Transistor is biased slightly into the “on” condition, usually specified as a constant quiescent drain current in FETs, with the DC gate bias adjusted to provide the design standing current. The optimal load is generally very close to the Class B value, and the achievable energy efficiency can be over 70%, theoretically. |
| C         | The Power transistor is biased below threshold, so under quiescent conditions, the transistor is switched off. This mode can be very energy efficient, but the gain is low, and the harmonic output and intermodulation distortion can be considerable.                                                                            |
| D, E, & F | Very efficient. No overlap between the voltage and current waveforms, so this amplifier converts the DC to RF with 100% efficiency. A resonator can be used to filter out the harmonics at the load.                                                                                                                               |
| D         | Form of switching amplifier, where the PA is driven by a pulsed waveform and the output current is switched rapidly between on and off states. Uses two transistors in a push-pull arrangement.                                                                                                                                    |
| E         | Both 2nd, and 3rd harmonics are tuned                                                                                                                                                                                                                                                                                              |
| F         | 2nd harmonic tuned                                                                                                                                                                                                                                                                                                                 |
| Inverse F | 2nd harmonic tuned                                                                                                                                                                                                                                                                                                                 |

The CATV industry primarily uses class A PA products because of high performance requirements, including low noise and high linearity, over a broad range of spectrum, typically between 5 and 1,218

MHz. While peak efficiencies for some of these classes may be high, they are attained only at or close to the maximum RF output power and will fall rapidly with OBO. To overcome this basic drawback, alternative PA architectures will be required.

#### 4.1. NLD Modeling

Perhaps the simplest way to represent PA NLD is to describe what it is not. A distortionless PA would have a linear transfer characteristic, which is achieved when the waveform of the PA's output voltage precisely duplicates that of its input. Only when a PA distorts does the output signal contain additional components, at frequencies differing from the frequencies of the input signal. The nature as well as the degree of distortion is dependent not only on the shape of the transfer characteristic of the PA, but also on the loading condition and operating point (bias) [12].

At higher powers, we have seen that the output power and gain deviate significantly from the linear relationship at small signals. This is the compression region of operation, and at a sufficiently high input drive, we will get no more power out of the PA -- at this point, we are at the saturated power level. In these regions of operation, the PA is very nonlinear. This compression behavior is also known as amplitude modulation to amplitude modulation (AM-AM) conversion: by modulating (changing) the input signal amplitude, we affect or modulate the amplitude of the output signal in a nonlinear fashion [12].

Typically, frequency-domain polynomial models will be used to model the AM-AM and amplitude modulation to phase modulation (AM-PM) characteristics of the PA. In general, frequency-domain models can describe the RF frequency response phenomena quite well but are unable to accommodate the memory effects associated with long time constants -- for example, bias line reactance and charge storage [10]. A Volterra series can be thought of as a Taylor series with memory; that is, a Taylor series defines not only at the present instant in time, but includes terms at previous instants, up to some specified delay [10].

Consider the baseband discrete model of the PA  $y(n) = f(x(n), x(n-1), \dots)$ , where  $x(n)$  is the input signal,  $y(n)$  is the output signal, and  $f(\cdot)$  is some nonlinear function. The simplest nonlinear PA model is the AM-AM/AM-PM model. Let the amplitude of the input signal be  $V_x = E[|x(n)|]$ , where  $E[\cdot]$  denotes short-term expectation or average. The output sample  $y(n)$ 's amplitude  $V_y = E[|y(n)|]$  and additional phase change  $\psi_y = E[\angle y(n)]$  depend on  $V_x$  in nonlinear ways as (6) and (7):

$$V_y = \frac{gV_x}{\left(1 + \frac{gV_x}{c}\right)^{\frac{1}{2\sigma}}} \quad (6)$$

$$\psi_y = \frac{\alpha V_x^p}{1 + \left(\frac{V_x}{\beta}\right)^q} \quad (7)$$

where  $g$  is the linear gain,  $\sigma$  is the smoothness factor, and  $c$  denotes the saturation magnitude of the PA. Typical examples of these parameters are  $g = 4.65$ ,  $\sigma = 0.81$ ,  $c = 0.58$ ,  $\alpha = 2560$ ,  $\beta = 0.114$ ,  $p = 2.4$ , and  $q = 2.3$ , which are used in the PA models regulated by Institute of Electrical and Electronics Engineers (IEEE) 803.11ad task group (TG) [11].

More accurate models should take into consideration the fact that nonlinearity leads to memory effects. In this case, Volterra series (8), are typically used to model PAs [11]. A general model is shown in [11] with up to Pth order nonlinearity and up to D step memory.

$$y(n) = \sum_{d=0}^D \sum_{k=1}^P b_{kd} x(n-d) |x(n-d)|^{k-1} \quad (8)$$

It can be shown that estimation of only odd-order nonlinearity (i.e. odd k) may be necessary for limited narrowband loading conditions and specific center frequencies, because even-order nonlinearity falls outside of the passband and will be filtered out by the receiver bandpass filters [11]. To illustrate this phenomenon, we can consider some simple examples where the input signal  $x(n)$  consists of a few (1-3) single frequency components only. Omitting the memory effects, if  $x(n)$  is a single frequency signal, i.e.,  $x(n) = V_0 \cos(a_0 + \phi)$ , where  $a_0 = 2\pi f_0 n$ . Then, using well-known trigonometric identities, the output signal can be written as

$$y(n) = k_1 V_0 \cos(a_0 + \phi + \psi_1) + \left( \frac{3}{4} k_3 V_0^3 + \frac{5}{8} k_5 V_0^5 \right) \cos(a_0 + \phi + \psi_3 + \psi_5) \quad (9)$$

$$+ \frac{1}{2} k_2 V_0^2 + \frac{3}{8} k_4 V_0^4 \quad (10)$$

$$+ \left( \frac{1}{2} k_2 V_0^2 + \frac{1}{2} k_4 V_0^4 \right) \cos(2a_0 + 2\phi + 2\psi_2 + 2\psi_4) \quad (11)$$

$$+ \left( \frac{1}{4} k_3 V_0^3 + \frac{5}{16} k_5 V_0^5 \right) \cos(3a_0 + 3\phi + 3\psi_3 + 3\psi_5) \quad (12)$$

$$+ \frac{1}{8} k_4 V_0^4 \cos(4a_0 + 4\phi + 4\psi_4) + \dots \quad (13)$$

$$+ \frac{1}{16} k_5 V_0^5 \cos(5a_0 + 5\phi + 5\psi_5) + \dots \quad (14)$$

$V_0 \cos(a_0 + \phi)$  and  $V_0 \sin(a_0 + \phi)$  are both sinusoidal voltages. Their waveforms are identical except for a 90° phase difference. The cosine form is used throughout this analysis because it results in simpler expressions [11]. The first line (9) is the in-band response with AM-AM/AM-PM nonlinear effects, the second line (10) is the direct current (DC) bias, and the remaining lines (11) through (14) include second through fifth order harmonics. At the receiving side, we may just have line (9) left, because all the other items will be canceled by bandpass filtering. Communication network filters, such as the root-raised cosine (RRC) filters, are typically implemented in two halves, one in the transmitter and the other in the receiver, so that overall, we get Nyquist rate sampling, and provide necessary impedance matching of the power transistor to its optimum load impedance of the network. Another critical filter function is in their use of controlling out-of-band emissions from sources including PA NLD, thus limiting the impact of distortion to within the operating band.

If  $x(n)$  consists of two frequencies, i.e.,  $x(n) = V_1 \cos(a_1 + \phi_1) + V_2 \cos(a_2 + \phi_2)$ , where  $a_i = 2\pi f_i n$ , then the inband response includes many more terms, such as the first order terms  $k_1 V_i \cos(a_i + \phi_i + \psi_i)$ , the third order terms  $k_3 (V_i^3 + V_i V_j^2) \cos(a_i + \phi_i + \psi_i)$ , the fifth order terms  $k_5 (V_i^5 + V_i V_j^4 + V_i^3 V_j^2) \cos(a_i + \phi_i + \psi_i)$ , for  $i, j \in \{1, 2\}$ . There are also intermodulation terms that consist of  $na_i \pm ma_j$  as long as they are within the passband of the bandpass filter, such as  $(V_i^2 V_j + V_i^2 V_j^3 + V_1^4 V_j) \cos(2a_i - a_j + 2\phi_i - \phi_j + 2\psi_i - \psi_j)$ . For some specific loading conditions, there may be many other higher order terms with frequencies  $na_i$ ,  $n(a_i \pm a_j)$ , or  $na_i + ma_j$ , that can not pass the passband filter. One of the important observations is that the contents that can pass the passband filter may consist of odd-order nonlinearity only for specific center frequency and narrowband conditions only.

If  $x(n)$  consists of three or more frequencies, we can have similar observations, albeit the expressions are more complex. Let the input signal  $x(n)$  be

$$x(n) = \sum_{i=1}^3 V_i \cos(a_i), a_i = 2\pi f_i n. \quad (15)$$

The second order component includes the DC component  $g_{2,0}(n)$ , the sum/difference of beat components  $g_{2,1}(n)$ , and the second-order harmonic components  $g_{2,2}(n)$ . Specifically,

$$k_2 x^2(n) = g_{2,0} + g_{2,1}(n) + g_{2,2}(n) \quad (16)$$

Where

$$g_{2,0}(n) = \sum_{i=1}^3 \frac{V_i^2}{2}, \quad (17)$$

$$g_{2,1}(n) = \sum_{i=1}^3 \sum_{j \neq 1} V_i V_j \cos(a_i \pm a_j), \quad (18)$$

$$g_{2,2}(n) = \sum_{i=1}^3 V_i^2 \frac{\cos(2a_i)}{2}. \quad (19)$$

The third order component includes the third-order harmonic components  $g_{3,1}(n)$ , the third intermodulation beat components  $g_{3,2}(n)$ , the triple beat components  $g_{3,3}(n)$ , the self-compression/expansion components  $g_{3,4}(n)$ , and the cross-compression/expansion components  $g_{3,5}(n)$ . Specifically,

$$k_3 x^3(n) = g_{3,1} + g_{3,2}(n) + g_{3,3}(n) + g_{3,4}(n) + g_{3,5}(n) \quad (20)$$

Where



$$g_{3,1}(n) = \sum_{i=1}^3 \frac{V_i^3}{4} \cos(3a_i), \quad (21)$$


$$g_{3,2}(n) = \sum_{i=1}^3 \sum_{j \neq 1} \frac{3V_i^2 V_j}{4} \cos(2a_i \pm a_j), \quad (22)$$

$$g_{3,3}(n) = \sum_{i=1}^3 \sum_{j \neq 1} \sum_{k \neq 1} \frac{3V_i V_j V_k}{2} \cos(a_i \pm a_j \pm a_k), \quad (23)$$

$$g_{3,4}(n) = \sum_{i=1}^3 \frac{3V_i^3}{4} \cos(a_i). \quad (24)$$

$$g_{3,5}(n) = \sum_{i=1}^3 \sum_{j \neq 1} \frac{3V_i V_j^2}{2} \cos(a_i). \quad (25)$$

The simulated output containing continuous wave signals (CWs) and NLD aligns with measurement in Figure 8, given coarse approximations for nonlinear gain coefficients and odd-order memory, based on model described in (8). A Rohde and Schwarz DOCSIS Cable Load Generator (CLGD) generated the CWs. The CWs propagated through a nonlinear power amplifier. The resultant nonlinear output spectrum was measured using a Keysight vector signal analyzer, model MXA, running in spectrum analysis mode.



**Figure 8 - 3 Tone NLD Measured vs. Simulated**


## 5. NLD Mitigation

The focus of this paper up to this point has been on the levers associated with PA efficiencies (a) near saturation operation, and more efficient implementations via (b) PA classes, which may include (c) biasing the PA to consume less power. All these approaches result in degraded NLD, which becomes even more challenging as modern-day orthogonal frequency division multiplexing (OFDM) signals, which have a higher peak-to-average power ratio (PAPR), become more ubiquitous in communication network payloads [15]. One of the major design goals modern systems is to make the communication systems more power efficient. This needs efficient PAs, which is unfortunately more challenging since OFDM has much higher PAPR and wider bandwidth [11].

We will next explore a portfolio of current methods focused on harvesting PA efficiency and in most cases include mitigating NLD. These methods have enabled network designers to push the network operating boundaries that were previously constrained by lower amounts of NLD. These methods will include PAPR Reduction, Transmitter DPD, and Receiver Post-Distorter Equalization.

### 5.1. Peak-to-Average-Power-Ratio (PAPR) Reduction

Figure 9 illustrates how PAPR can interact with PA characteristics. OFDM signals can have a high degree of PAPR and push the PA operation into saturation at maximum signal amplitudes, also illustrated in Figure 9. PAPR reduction methods work to minimize the signal's PAPR, so that operation at PA saturation can be avoided, like OBO. There are multiple methods available to achieve PAPR reduction, as we will see, each with their own tradeoffs in benefit and cost.



**Figure 9 - PA Characteristics with PAPR and DPD**

PAPR reduction was discussed during the deliberations of the DOCSIS 3.1 standard development, which introduced OFDM signaling into its portfolio of physical layer technology (PHY) [7]. The concern was that the increased PAPR of OFDM exceeded that of previous generation DOCSIS 3.0, SC-QAM on a 6 MHz bandwidth basis. However, the standards group ultimately decided against using PAPR reduction, primarily because the impact to CATV PAs would be negligible, given their broadband nature, which is an aggregate of over one hundred 6 MHz channels, up to approximately 1 GHz of bandwidth, total. Through lab measurement, these conclusions have been validated, but actual field results at scale have yet to be made [15].

High data rates have led to complex modulation schema, and ultimately higher spectral efficiency. One consequence of using spectrally efficient modulation schemes is that the dynamic range of the signal may be quite high [10]. This is generally measured in terms of the ratio of the peak signal power to the average power of the modulated signal, or PAPR [10]. While a large PAPR is not such a problem for signal transmission, it can have an impact on the efficiency of the transmitter PA [10], as illustrated in Figure 9, where the signal peaks can push the PA into saturation if the signal OBO isn't sufficiently below the PA compression point. Reducing the PAPR of the input signal using digital signal processing allows the PA to be operated at a higher efficiency, and it also reduces the dynamic range needed to represent the input signal digitally [10]. In addition, the reduced PAPR will often reduce the complexity of the linearization approach needed to compensate for PA nonlinearities [10].

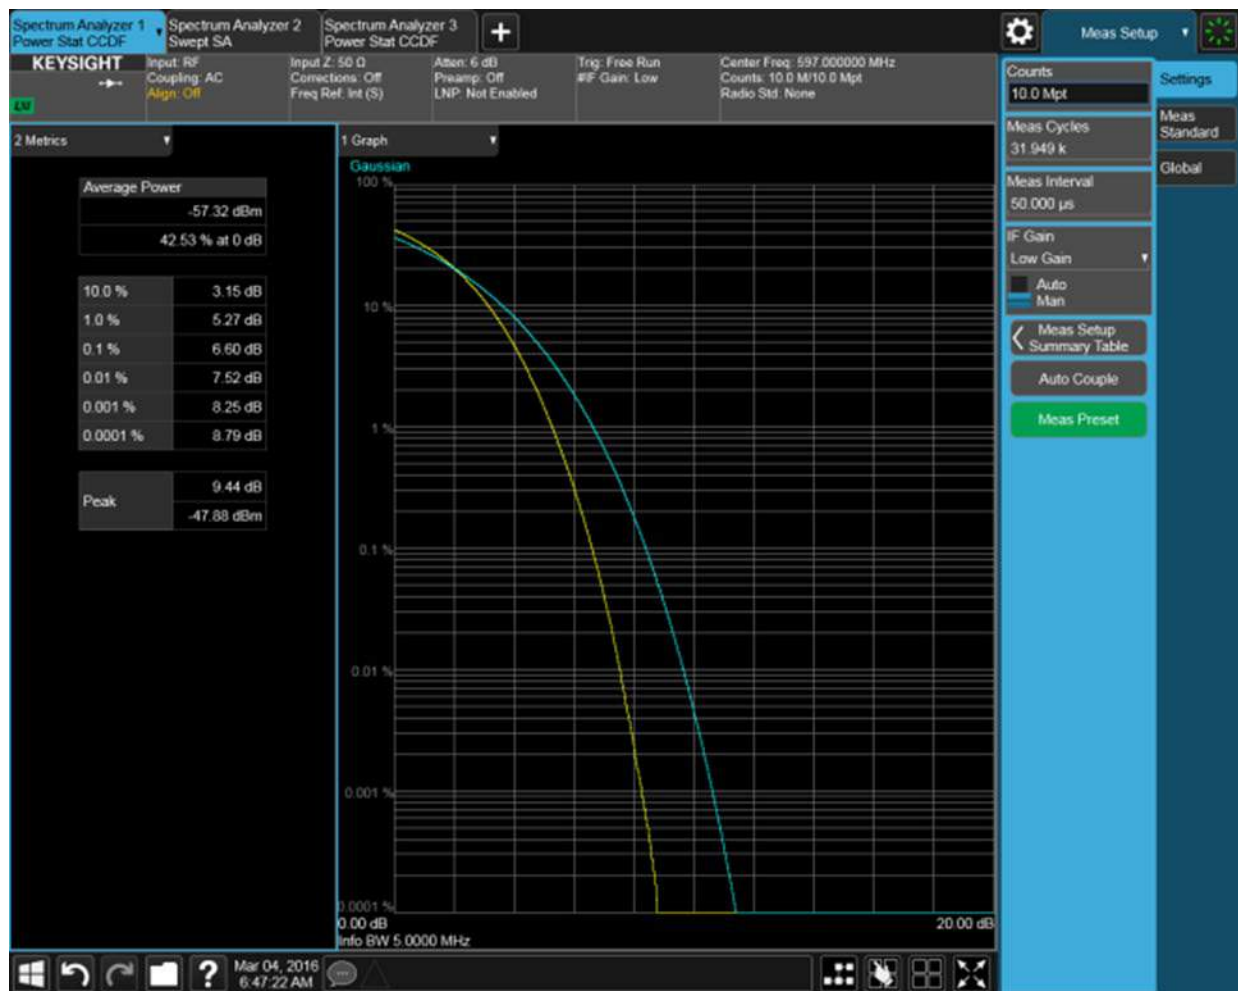
The PAPR of a signal  $x(n)$  is illustrated in (26), where  $E[\cdot]$  denotes short-term expectation or average [10].

$$PAPR = \frac{\max(|x|^2)}{E[|x|^2]} \quad (26)$$

(27) illustrates how PAPR is based on the statistics of the signal rather than the absolute peak, while the practical peak is the level,  $L$ , at which the signal magnitude has a  $10^{-4}$  probability,  $P$ , of exceeding.

$$P\{|x|^2 > L\} = 10^{-4} \quad (27)$$

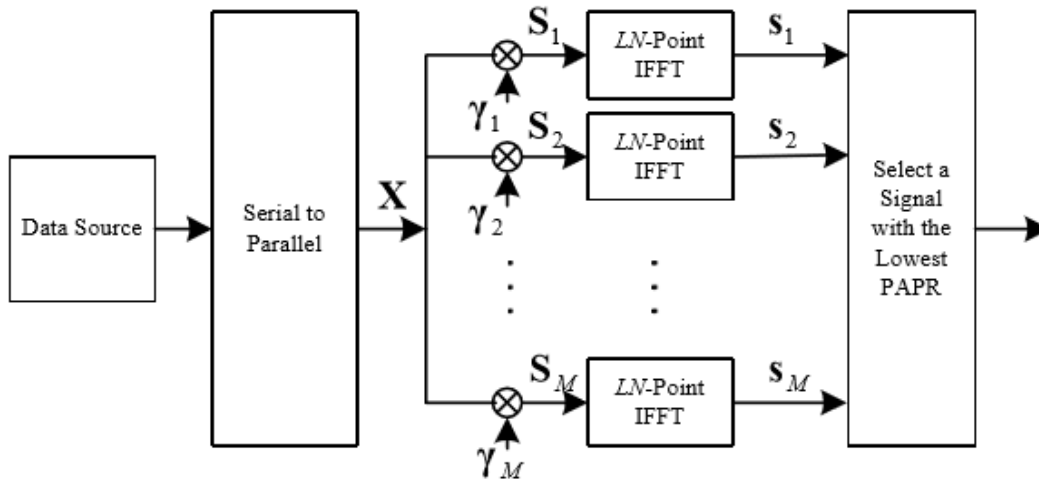
The complementary cumulative density function (CCDF) of  $|x|^2$ , is a useful description of the signal statistics, often compared to that of a Gaussian waveform because the statistics of multi-carrier signals used in many of today's communications networks tend to approach that of a Gaussian [10]. Figure 10 is a Keysight PXA PAPR measurement for a 6 MHz, DOCSIS downstream SC-QAM signal whose PAPR (yellow line) is approximately 9.44 dB based on a 10 Mpt (million point) sample period [15]. Changing from DOCSIS SC-QAM to an OFDM-based PHY and limiting measurements to 6 MHz bandwidth, shows that PAPR will increase by approximately 0.96 dB, using M=256 QAM constellation levels for both signal types. However, much broader bandwidth comparisons over 192 MHz revealed that OFDM PAPR was lower than SC-QAM, by about 0.52 dB [15].



**Figure 10 - DOCSIS Downstream SC-QAM PAPR Measurement**


Many techniques have been developed for PAPR reduction, such as signal clipping, peak cancellation, and error waveform subtraction (noise shaping) [10]. These clip-and-filter approaches clip peaks exceeding a specified level and filter the waveform to remove out-of-band distortion [10]. Clip-and-filter and peak windowing are the easiest Crest Factor Reduction (CFR) methods to implement, making them the most likely to be used [10].

For OFDM-based formats, CFR is often achieved using redundant coding or the transmission of auxiliary information, both of which reduce data throughput of the system. Pilot tones and unmodulated subcarriers can be exploited to reduce PAPR with some special pre-coding techniques [16]. The selected mapping approach (SLM) from [16] provides good PAPR reduction performance, but may suffer from high computational complexity from using a bank of inverse fast fourier transforms (IFFTs), illustrated in Figure 11.



**Figure 11 - Traditional Selected Mapping, PAPR Reduction via Pre-Coding**

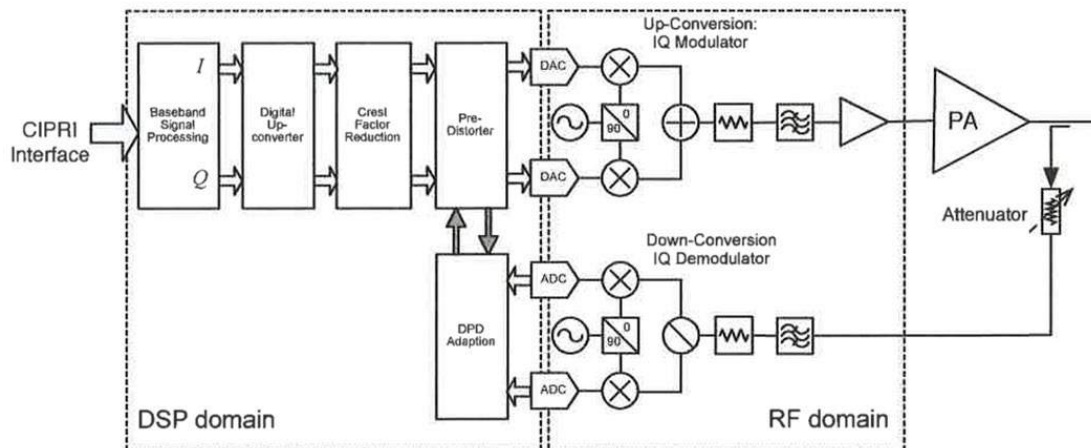
The new SLM approach from [16], shown in Figure 12, replaces IFFTs with new kinds of conversions, resulting in much lower complexity. Multiple IFFTs are replaced with transformation matrix,  $T_r$ , to produce candidate signals [16]. PAPR reduction for this approach is almost as good as traditional SLM approaches but uses much less processing overhead [16].



**Figure 12 - Transformation Based Pre-Coding for SLM PAPR Reduction**

CFR techniques are very appropriate for many wireless communications applications, such as cellular wireless handset PA applications, where Class AB operations are typically used [10]. CFR is a form of digital signal processing applied to the digital signal  $x(n)$ , used in combination with DPD to reduce the requirements on the RF PA within the transmit chain [10]. Although the distortionless PAPR reduction methods decrease the deleterious effect of nonlinear distortions, their effectiveness in improving the system performance is limited, since the main problem of the limited dynamic range of the PA remains unsolved [17]. Therefore, CFR often precedes DPD, per Figure 13. When CFR is used in conjunction

with DPD techniques, the expansive nature of the DPD function operates on the crest factor reduced signal, so that the resulting pre-distorted signal that enters the PA does not have an excessive PAPR, and the PA can still be driven hard to operate at its highest efficiency [10]. This behavior is also illustrated in Figure 9.




**Figure 13 - High-Level Transmitter with CFR and Adaptive DPD**

CFR assumes that reducing the peak of the signal is beneficial to the performance of the digital transmitter, essentially allowing the linearized PA to operate at a higher efficiency, while meeting the linearity requirements of the modulation format used [10]. CFR has drawbacks which degrade the system performance, either by increasing in-band degradation or reducing data throughput, but the overall system can be optimized to meet the performance specifications at higher PA efficiencies [10]. When applied properly, CFR allows the PA to operate more efficiently, thereby improving the performance of the transmitter [10].

### 5.1. Transmitter Digital Pre-Distortion (DPD)

One of the dominating practices in the cellular and satellite industries today is to insert a DPD circuit before the PA, which distorts the signals appropriately to compensate for the nonlinear PA response [11]. The distortion is added in such a way as to cancel the inherent nonlinearity of the PA, so that its output is a linear replica of the original input signal. DPD has been applied widely in many modern transmitters [2]. DPD can lead to the use of more efficient and cost-effective PAs [11] and is being considered for future generations of FDX RPD node hardware, where power consumption thresholds are already being encroached upon, while new capability is being added to increase its capacity as efficiently as possible.

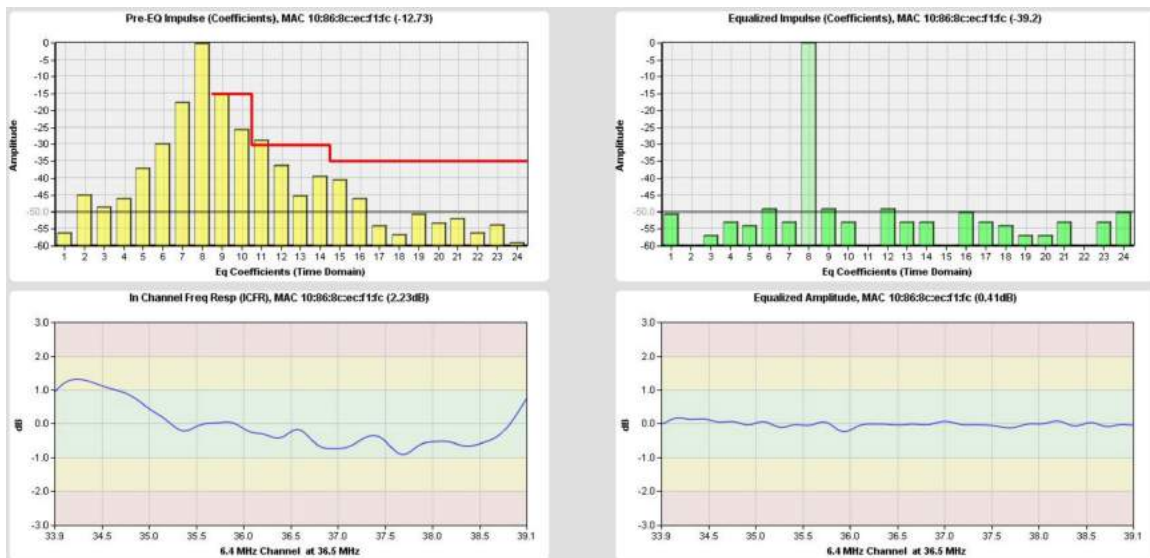


**Figure 14 - Nonlinearized PA vs. Linearized PA via Digital Pre-Distortion**

We will now focus on current approaches available to mitigate the effects of NLD, which are based on equalization, or creating an anti-phase version of NLD, as in the bottom part of Figure 14, which, when added to the signal from which it was derived, can negate the nonlinear effects of the PA. This essentially linearizes the PA's behavior.

In Figure 14, equalization is being performed prior to the signal's exposure to NLD, or pre-equalization. Pre-equalization of NLD is typically accomplished using DPD technology. DPD is analogous to DOCSIS Transmit Pre-Equalization, and the key difference is in the type of distortion that gets mitigated. DOCSIS Transmit Pre-Equalization mitigates linear distortion (LD), or plant echoes and other filter effects including amplitude roll-off, and group delay variation [7]. The key similarity here is that both signal processing techniques are applied as a signal bias, prior to the signal's impairment exposure, impairments being NLD for DPD, and LD for transmit pre-equalization.





**Figure 15 - CM Upstream Transmit Pre-Equalization and Post-Equalization Functions**


Equalization functions can also be applied to the signal after impairment exposure, and this is known as post-distorter equalization for NLD and post-equalization for LD. In DOCSIS upstream communications, both Transmit Pre-Equalization and post-equalization functions work collaboratively to mitigate the effects of LD [7]. Figure 15 illustrates these two functions together compensating for the communication channel that exists between, say, my CM, and the CMTS that serves multiple neighboring towns, including my own. CMTSs may connect to hundreds of nodes, as discussed in earlier sections, and ultimately connecting to thousands of CMs. The CMTS's primary function is that of a router, facilitating communication between the local area network (LAN), comprised of many CMs, including mine, and the wide area network (WAN) or internet.

The collaboration between the CM and CMTS on how to equalize upstream LD is specified in DOCSIS, where transmit pre-equalization is a CM function and its DOCSIS 2.0, 24 symbol-spaced coefficients, shown in the upper left side of Figure 15 in yellow, are provided to it by the CMTS [7]. The CMTS post-equalization function is shown in the upper right side, in green. While performing its own post equalization function, the CMTS periodically sends a set of equalizer coefficients to the CM, via station-maintenance messages, with instructions to either overwrite or convolve the CM's current set of equalizer coefficients with the new coefficients.

An effect of this collaboration is to perform most of the channel equalization at the CM's transmit pre-equalization function, leaving only minor corrections in the post equalization function, at the CMTS. The equalizer's coefficients are colored yellow because of the intensity of correction, or, in other words, the variation of its amplitude frequency response, illustrated in the lower left chart, which exceeds  $\pm 1$  dB peak-to-peak. Figure 15 illustrates how the CM is compensating for most of the channel's LD. Overall, LD equalization at both ends of the communication link has proven to be very robust and reliable in the CATV industry, but it does so with additional signal processing and overhead.

Adaptive pre-distortion techniques include an observation path that samples the output from the power amplifier and feeds the signal back to the pre-distorter to estimate system nonlinearity [10]. Additional components include analog-to-digital converters (ADCs) which are used to convert the observation signal back to baseband and digitize for digital domain operations. Components in this path will have better linearity than the desired performance of the transmitter to avoid introducing distortion at the PA output

arising from the observation process [10]. Input and feedback signals must be synchronized, accounting for the group delay of the PA, which tends to dominate the loop [10].




**Figure 16 - Effects of Sample Rate on a Pre-Distorted Signal**

Pre-distorters reconstruct the NLD in the analog domain, for input into the PA chain in the transmitter, which needs to be sampled at a higher rate to accommodate for the increased bandwidth. Therefore upsampling is required, otherwise under sampling the NLD could lead to aliasing in Figure 16 [10]. The output signal from the pre-distorter (b) has a much wider bandwidth than the input signal (a). Under-sampling of the input will lead to aliasing of the output signal (c), and the clean DPD. Typical oversampling by 5x will represent 5th order NLD [10]. Two main approaches have been adopted in the pre-distorter block:

- (1) Look-up tables (LUTs) to provide a map relating the desired pre-distorter output to the input voltage

- (2) Nonlinear basis functions to describe nonlinear pre-distortion function, requiring generation of the nonlinear functions and the multiplication of the basis functions by the input voltage [10].

Many PA models can describe the nonlinear pre-distortion function, including the Volterra series and Wiener model [19]. PA models will need to consider that PA characteristics do not change rapidly with time; changes in PA characteristics are often attributable to temperature drift, aging, etc., which have long time constants [19]. The cause of memory effects can be electrical or electrothermal [19]. Higher output power PA operation, such as those used in wireless base-stations, exhibit memory effects [19]. Having memory means that the output of the PA is not only a function of the current input, but also a function of past inputs and outputs [18]. Memory effects in the power amplifier limit the performance of DPD for wideband signals [18].



**Figure 17 - Memory and Memoryless DPD Results**

An issue with DPD is that it is compensating for the PA in the transmitter only, which, in most practical networks today, represents a fraction of the nonlinearity present in the communication channel. PAs are at least used in both the transmitter and receiver. Multiple PAs could be used in between the transmitter and receiver to extend the reach in many communication networks, such as CATV networks. Passive components are used to distribute signals to many receivers, and may also contain nonlinear components, such as inductors, that contribute to the channel's aggregate nonlinear NLD.

AI methods for DPD, using neural networks, illustrate the potential for improvement for NLD performance over traditional methods [20] [21]. However, incorporating additional NLD compensation

into the transmitter to account for the rest of the communication network chain, like similar closed-loop equalization systems, could lead to DPD reaching its full potential.

Implementations of DPD should be architected to compensate for stronger nonlinearity, much like the closed-loop LD equalization strategy described in the beginning of this section, for DOCSIS upstream signals can compensate for appreciably high LD. Because of bidirectional network connectivity, both the CMTS and CM can collaborate on the estimation of the total NLD present within the communication channel, and send coefficients via the downstream communication path, with instructions to either convolve or overwrite coefficients describing the estimated path NLD that will account for all the nonlinear elements within the communication network chain (transmitter, receiver, amplifiers, and nonlinear passive components). For this strategy to work, receivers must be capable of mitigating NLD. We will next review the current solutions available in receiver-based post distortion cancellation.

## **5.1. Receiver Post-Distorter Equalization**

Another strategy is to mitigate the nonlinear PA distortions at the receivers via post-distorter equalization [22] [23] [24]. The solution presented in [25] develops a Bayesian signal detection algorithm, based on the nonlinear response of the PAs. However, this documented approach applies to a simple “AM-AM, AM-PM” nonlinear PA model only.

The authors of [22] propose a symbol-based equalizer, with nonlinear distortion cancellation for the forward link as an addition to the standard linear equalizer at the receiver, suitably adapted to incorporate specific channel functions in the forward link, including input multiplexing (IMUX) and output multiplexing (OMUX) filtering and the traveling wave tube amplifier (TWTA), using the memory polynomial model. The proposed setup is compared with current mitigation approaches, yielding significant efficiency gains [22]. The nonlinearity in the satellite channel is introduced primarily by the TWTA [22].

The objective of [22] was to compare the performance of the proposed equalizer with the state-of-the-art dynamic data pre-distortion, and to show that the best system performance is achieved when both pre-distortion at the transmitter and decision-directed equalization at the receiver are applied. In addition, the performance of a simple maximum likelihood (ML) demodulator in the detector in the cancellation loop is compared against a low density parity check (LDPC) decoder, to show the robustness of the proposed equalizer to decision errors [22]. Overall system complexity of the proposed nonlinear equalizer is argued to be less than current linear equalization approach, due to less-frequent updates [22]. As discussed earlier in the section covering DPD, similar transmitter and receiver equalization loops have proven to be robust against LD for DOCSIS upstream [7].

References [23] [24] and [25] approach receiver post NLD equalization via clustering methods, which is a subset of broader suite of AI models. [23] leverages SVMs, while [24] uses a radial basis function (RBF) network and [25] is Bayesian. ANNs have attracted researchers in the field of PA modeling, due to its successful implementation in pattern recognition, signal processing, system identification, and control [26].

## **6. Severe NLD Mitigation**

One of the major design goals for modern systems is to make the communication systems more power efficient. This needs efficient PAs, which is unfortunately more challenging, since many modern PHY include OFDM, which has much higher PAPR and wider bandwidth [11].


Existing nonlinear PA mitigation strategies -- DPD, PAPR reduction, and receiver-based equalization, discussed in Section 5 -- may not be sufficient enough when considered individually. We can reduce PAPR to some extent only. DPD is too complex and costly for small and low-cost devices. Existing DPD and equalization techniques have moderate nonlinear distortion compensation capabilities, because they have been designed to cancel internal transmitter NLD only.

There is a larger system that must be considered, consisting of a transmitter, receiver, and nonlinear active/passive components in between them, all contributing to the overall NLD observed at the receiver. In some cases, the NLD can be quite severe, due to required higher-efficiency modes of operation. Supporting collaboration between the transmitter and receiver, like the LD cancellation systems used for DOCSIS, may be the key to achieving optimal network efficiency, while minimizing NLD. Severe nonlinearity estimation and mitigation is a requirement for the receiver, representing information that could then be shared with the transmitter through bidirectional communication. Ultimately, this lessens the burden at the receiver location and the NLD equalization system overall.

In this section, a system for cancelling severe NLD will be proposed which develops nonlinear equalizers that exploit both deep neural networks (DNNs) and Volterra series models to mitigate PA nonlinear distortions. The DNN equalizer architecture consists of multiple one-dimension convolutional layers. The input features are designed according to the Volterra series model of nonlinear PAs. This enables the DNN equalizer to mitigate nonlinear PA distortions more effectively, while avoiding over-fitting under conditions of limited training data. Experiments are conducted with real measurement data obtained from a highly nonlinear RFMD RF2317 Linear CATV Amplifier [11]. The results will demonstrate that the proposed DNN equalizer has superior performance over conventional equalization approaches, a necessary tool for more collaborative NLD mitigation strategy that could make more efficient network components potentially realizable.


## 6.1. Enhanced Equalization – Integrating Volterra Series and DNNs

Reference [11] proposes the use of DNNs to implement the nonlinear equalizer in the receiver, which can mitigate the nonlinear effects of the received signals, due to not only PAs but also nonlinear channels and propagations. The architecture of the DNN equalizer is shown in Figure 18.



**Figure 18 - Block Diagram of DNN Equalizer**

Different from [28], [11] used multi-layer convolutional neural networks (CNNs). Different from conventional neural network predistorters proposed in [29], [11] used neural networks as equalizers at the receivers. Different from conventional neural network equalizers such as those proposed in [30] [31], in a DNN equalizer [11], a CNN is used and the input features in  $X$  are not only the linear delayed samples  $r(n)$ . Rather, Volterra series models are used to create input features.



**Figure 19 - System Block Diagram with Nonlinear Power Amplifier and Deep Neural Network Equalizer**

To simplify presentation, according to the previous section, [11] assumes that the linear channel  $H$  has already been equalized by a linear equalizer of Figure 19, whose output signal is  $r(n)$ . According to Volterra series representation of nonlinear functions, the input-output response of the nonlinear equalizer can be written as

$$z(n) = \sum_{k=1}^P \sum_{d_1=0}^D \cdots \sum_{d_k=0}^D f_{d_1, \dots, d_k} \prod_{i=1}^k r(n - d_i) \quad (28)$$

One of the major problems is that the number of coefficients  $f_{d_1, \dots, d_k}$  increases exponentially with the increase of memory length  $D$  and nonlinearity order  $P$ . There are many different ways to develop more efficient Volterra series representations with reduced numbers of coefficients. For example, in [32], the authors exploit the fact that higher-order terms do not contribute significantly to the memory effects of PAs to reduce the memory depth  $d$  when the nonlinearity order  $k$  increases. This technique can drastically reduce the total number of coefficients. In [33] [34] [35], the authors developed the dynamic deviation model to reduce the full Volterra series model (28) to the following simplified one

$$\begin{aligned} z(n) &= z_s(n) + z_d(n) \\ &= \sum_{k=1}^P f_{k,0} r^k(n) + \sum_{k=1}^P \sum_{j=1}^k r^{k-j}(n) \sum_{d_1=0}^D \cdots \sum_{d_j=d_{j-1}}^D f_{k,j} \prod_{i=1}^j r(n - d_i) \end{aligned} \quad (29)$$

where  $z_s(n)$  is the static term, and  $z_d(n)$  is the dynamic term that includes all the memory effects. We can see that the total number of coefficients can be much reduced by controlling the dynamic order  $j$  which is a selectable parameter.

[11] constructs the input features of the DNN based on the model (29). Corresponding to the static term  $z_s(n)$ , [11] changes it to

$$\hat{z}_s(n) = \sum_{1 \leq k \leq P} f_{k,0} r(n) |r(n)|^{k-1} \quad (30)$$

The reason that (30) changes  $r^k(n)$  to  $r(n)|r(n)|^{k-1}$  is that only the signal frequency with the valid passband is of interest. This means that input feature vector  $X$  should include terms  $r(n)|r(n)|^{k-1}$ .

Similarly, corresponding to the dynamic term  $z_d(n)$ , we need to supply  $r^{k-j}(n) \prod_{i=1}^j r(n - d_i)$  in the features where half of the terms  $r(n)$  and  $r(n - d_i)$  should be conjugated. For simplicity, in DNN equalizer used in [11], the vector  $X$  includes  $r(n - q)|r(n - q)|^{k-1}$  for some  $q$  and  $k$ .

By applying Volterra series components directly as features of the input  $X$ , the DNN can develop more complex nonlinear functions with fewer hidden layers and fewer neurons. This will also make the training procedure converge much faster, with much less training data.

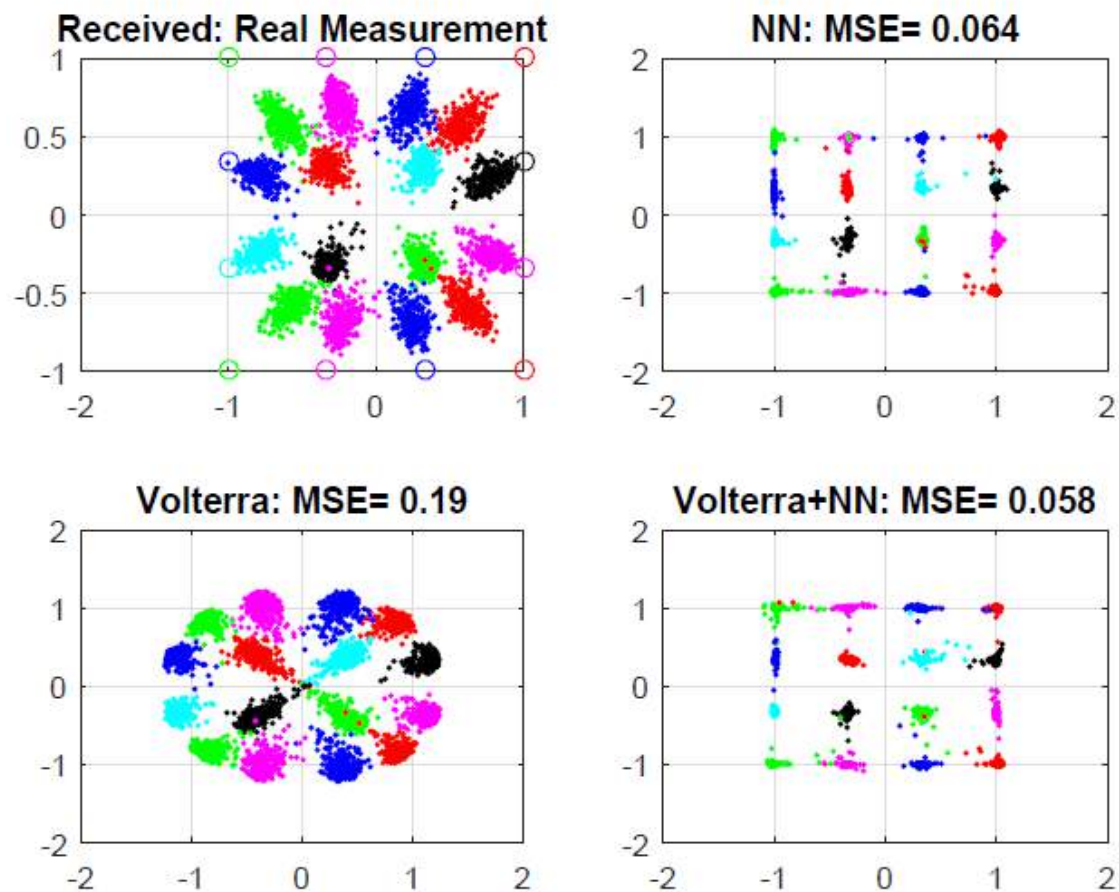
In Figure 18, the input  $X$  is a tensor formed by the real and imaginary parts of  $r(n - q)|r(n - q)|^{k-1}$  with appropriate number of delays  $q$  and nonlinearities  $k$ . There are three one-dimension convolutional layers, each with 20 or 10 feature maps. After a drop-out layer for regularization, this is followed by a fully connected layer with 20 neurons. Finally, there is a fully connected layer to form the output tensor  $Y$  which has two dimensions. The output  $Y$  is used to construct the complex  $z(n)$ , where  $z(n) = \hat{x}(n - d)$  for some appropriate delay  $d$ . All the convolutional layers and the first fully connected layer use the sigmoid activation function, while the output layer uses the linear activation function. [11] uses the mean square error loss function  $L_{loss} = E[|x(n - d) - z(n)|^2]$ , where  $z(n)$  is replaced by  $Y$  and  $x(n - d)$  is replaced by training data labels.

Measurement signals were obtained from an implementation of a RFMD RF2317 PA used in the cable industry, which are typically dominated by 3rd order nonlinearities. Various levels of nonlinear distortion, in terms of dBc, were generated by adjusting the PA RF input levels [11].

For the Volterra equalizer, reference [11] approximated the response of the nonlinear equalizer with delays including 8 pre- and post- main taps, and with nonlinearity including even and odd order nonlinearity up to the 5th order. To determine the values of the Volterra coefficients, we transmitted  $N = 4,096$  training symbols through the PA and then collected the noisy received samples  $r(n)$ .

For conventional time-delay NN equalizer, [11] applied a feedforward neural network with 80-dimensional input vector  $X$  and 5 fully connected hidden layers with 20, 20, 10, 10, 10 neurons, respectively.






**Figure 20 - Constellation of 16-QAM Non-Equalized vs. Equalized**

Figure 20 shows the constellation of 16-QAM equalization over the real PA. The corresponding SER were 0.0067, 0.0027, 0.00025 respectively. Nonlinear filtering of both a DNN and DNN with Volterra input features show superior equalization over just Volterra filtering alone. Further, it can be seen that the proposed Volterra+NN scheme has the best performance.





**Figure 21 - Comparing Three Equalization Methods for 16-QAM under Various NLD Levels**

Figure 21 provides MSE measurements for 16-QAM under various nonlinear distortion levels, dBc. For each 1 dB increase in NLD, the resultant MSE is shown for the “Measured”, “Volterra”, “NN”, and the proposed “Volterra+NN” cases. MSE reduction diminishes appreciably as modulation order increases from QPSK to 64-QAM, but small improvements in MSE have been observed to lead to appreciable SER improvement, especially for more complex modulation orders. Unfortunately, the 4,096 symbol sample size limited measurements to a minimum measurable 0.000244 SER, which represents 1 symbol error out of 4,096 symbols.

**Table 3 - Comparing MSE/SER Improvement % for the Three Equalization Methods**

|                | Volterra   |            | NN          |            | Volterra+NN |            |
|----------------|------------|------------|-------------|------------|-------------|------------|
|                | MSE        | SER        | MSE         | SER        | MSE         | SER        |
| <b>64-QAM</b>  | <b>16%</b> | <b>26%</b> | <b>10%</b>  | <b>25%</b> | <b>42%</b>  | <b>44%</b> |
| <b>16-QAM</b>  | <b>41%</b> | <b>2%</b>  | <b>35%</b>  | <b>6%</b>  | <b>85%</b>  | <b>28%</b> |
| <b>QPSK</b>    | <b>57%</b> | <b>0%</b>  | <b>100%</b> | <b>0%</b>  | <b>100%</b> | <b>0%</b>  |
| <b>AVERAGE</b> | <b>38%</b> | <b>9%</b>  | <b>48%</b>  | <b>10%</b> | <b>76%</b>  | <b>24%</b> |

Table 3 summarizes equalization performance, which shows the average percent reduction/improvement in MSE and SER from NLD-impaired data for multiple modulation orders.

## 7. Conclusion

The enhanced capacity associated with FDX is increasing the need for higher efficiency PAs. A familiar scenario also playing out in the cellular, Wi-Fi, and satellite industries, where similar capacity-enhancing needs are increasing the needs for more efficient PAs. PA efficiencies can be realized by (a) near saturation operation, and more efficient implementations via (b) PA classes, which may include (c) biasing the PA to consume less power. All these approaches result in degraded NLD.

NLD mitigation techniques, like DPD, can lead to the use of more efficient and cost-effective PAs [11], and are being considered for future generations of FDX RPD node hardware -- where power consumption thresholds are already being encroached upon, while new capabilities are being added to increase their capacity as efficiently as possible. However, an issue with DPD is that it is compensating for the PA in the transmitter only, which, in most practical networks today, represents a fraction of the nonlinearity present in the communication channel. Incorporating additional NLD compensation into the transmitter, to account for the rest of the communication network chain, like similar closed-loop equalization systems discussed in this paper, could lead to DPD reaching its full potential.

Bidirectional network connectivity between the CMTS and CM can enable the convergence to an estimate of NLD present within the communication channel, in either direction, and send coefficients via the downstream communication paths, with instructions to either convolve or overwrite coefficients describing the estimated path NLD that will account for all the nonlinear elements within the communication network chain (transmitter, receiver, amplifiers, and nonlinear passive components).

However, for this strategy to work, receivers must be capable of mitigating severe NLD. Supporting collaboration between the transmitter and receiver, like the LD cancellation systems used for DOCSIS, may be the key to achieving optimal network efficiency, while minimizing NLD.

Thus, more aggressive NLD cancellation methods may be accomplished by advanced DNN approaches, such as incorporating input features derived from Volterra series models. Results from [11] demonstrate that the proposed DNN equalizer has superior performance over conventional equalization approaches, and is a necessary tool for more collaborative NLD mitigation strategy. Ultimately, this could make more efficient network components realizable.

## Abbreviations

|         |                                              |
|---------|----------------------------------------------|
| ACLR    | adjacent channel leakage ratio               |
| ACPR    | adjacent channel power ratio                 |
| ADC     | analog-to-digital converter                  |
| AI      | artificial intelligence                      |
| ANN     | artificial neural networks                   |
| AM-AM   | amplitude modulation to amplitude modulation |
| AM-PM   | amplitude modulation to phase modulation     |
| Auto ML | automated machine learning                   |
| CAGR    | compound annual growth rate                  |
| CATV    | cable television                             |
| CCDF    | complementary cumulative density function    |
| CFR     | crest factor reduction                       |
| CIN     | composite-intermodulation-noise              |
| CLGD    | DOCSIS Cable Load Generator                  |

|        |                                                   |
|--------|---------------------------------------------------|
| CM     | cable modem                                       |
| CMTS   | cable modem termination system                    |
| CNN    | convolutional neural network                      |
| CPE    | customer premise equipment                        |
| CW     | continuous wave                                   |
| dB     | decibel                                           |
| dBc    | decibels relative to carrier power                |
| dBmV   | decibel-millivolts                                |
| DC     | direct current                                    |
| DNN    | deep neural networks                              |
| DOCSIS | data over coax system interface specifications    |
| DPD    | digital pre-distortion                            |
| DTA    | digital terminal adapter                          |
| EC     | echo cancellation                                 |
| EOL    | end-of-line                                       |
| EVM    | error vector magnitude                            |
| FEC    | forward error correction                          |
| FDX    | full duplex DOCSIS                                |
| Hz     | hertz                                             |
| IEEE   | Institute of Electrical and Electronics Engineers |
| IFFT   | inverse fast fourier transform                    |
| IMD    | intermodulation distortion                        |
| IMUX   | input multiplex                                   |
| IoT    | internet of things                                |
| ISI    | inter-symbol interference                         |
| LAN    | local area network                                |
| LD     | linear distortion                                 |
| LDPC   | low density parity check                          |
| LMS    | least mean squares                                |
| LUT    | look-up table                                     |
| MER    | modulation error ratio                            |
| MHz    | mega-hertz                                        |
| ML     | maximum likelihood                                |
| Mpt    | million point                                     |
| mV     | milli-volt                                        |
| NF     | noise figure                                      |
| NLD    | nonlinear distortion                              |
| OBO    | output-power-back-off                             |
| OFDM   | orthogonal frequency division multiplexing        |
| OMUX   | output multiplex                                  |
| PA     | power amplifier                                   |
| PAPR   | peak-to-average-power-ratio                       |
| PHY    | physical layer                                    |
| RBF    | radial basis function                             |
| RF     | radio frequency                                   |
| RPD    | remote PHY device                                 |
| RRC    | root-raised cosine                                |
| QPSK   | quadrature phase shift keying                     |

|                  |                                                |
|------------------|------------------------------------------------|
| SC-QAM           | single carrier quadrature amplitude modulation |
| SLA              | service level agreement                        |
| SLM              | selected mapping approach                      |
| SNR              | signal-to-noise ratio                          |
| SNR <sub>s</sub> | system signal-to-noise ratio                   |
| STB              | set-top box                                    |
| SVM              | support vector machines                        |
| TG               | task group                                     |
| TWTA             | traveling wave tube amplifier                  |
| VSA              | vector signal analyzer                         |
| WAN              | wide area network                              |

## Bibliography & References

- [1] S. Haykin, Neural Networks: A Comprehensive Foundation. Prentice-Hall, Inc., 1999
- [2] P. Winston, Artificial Intelligence, <https://podcasts.apple.com/us/podcast/artificial-intelligence/id765641080>, MITOPENCOURSEWARE, 2013
- [3] OpenAI, Solving Rubik's Cube with a Robot Hand, <https://openai.com/blog/solving-rubiks-cube/>, 2019
- [4] edureka!, Top 12 Artificial Intelligence Tools & Frameworks you need to know, <https://www.edureka.co/blog/top-12-artificial-intelligence-tools/#scikit>
- [5] V. Fedak, Top 10 Most Popular AI Models, <https://dzone.com/articles/top-10-most-popular-ai-models>, 2018
- [6] R. Thompson, C. Moore, J. Moran, R. Howald, Optimizing Upstream Throughput Using Equalization Coefficient Analysis, <https://www.nctatechnicalpapers.com/Paper/2009/2009-optimizing-upstream-throughput-using-equalization-coefficient-analysis>, 2009
- [7] Cable Television Laboratories, Inc., Data Over Cable System Interface Specifications, DOCSIS© 1.1, Radio Frequency Interface Specification, <https://www.cablelabs.com/specifications/radio-frequency-interface-specification>, September 7<sup>th</sup>, 2005
- DOCSIS© 2.0, Radio Frequency Interface Specification, <https://www.cablelabs.com/specifications/radio-frequency-interface-specification-2>, April 22<sup>nd</sup>, 2009
- DOCSIS© 3.0, Physical Layer Specification, <https://specification-search.cablelabs.com/CM-SP-PHYv3.0>, December 7<sup>th</sup>, 2017
- DOCSIS© 3.1, Physical Layer Specification, <https://specification-search.cablelabs.com/CM-SP-PHYv3.1>, September 17<sup>th</sup>, 2019
- DOCSIS© 4.0, Physical Layer Specification, <https://www.cablelabs.com/specifications/CM-SP-PHYv4.0>, April 29<sup>th</sup>, 2020

- [8] Charles Warren, “How might we, three words that make design better”, <https://hbr.org/2012/09/the-secret-phrase-top-innovato>, <https://www.youtube.com/watch?v=mTpa-bJiMp4>
- [10] Cisco, “GS7000 1218MHz Fiber Deep Intelligent Node Data Sheet”, <https://www.cisco.com/c/en/us/products/collateral/video/gs7000-node/datasheet-c78-740828.html>, January 9, 2019
- [11] R. J. Thompson and X. Li, “Integrating Volterra Series Model and Deep Neural Networks to Equalize Nonlinear Power Amplifiers,” in 2019 53rd Annual Conference on Information Sciences and Systems (CISS). IEEE, 2019, pp. 1–6.
- [12] K. Simons, Technical Handbook for CATV Systems, 3rd Edition. Jerrod Publication No. 436-001-01, 1968.
- [13] W. Ciciora, J. Farmer, D. Large, Modern Cable Television Technology; Voice, Video, and Data Communications. Morgan Kaufmann Publishers, Inc. 1999.
- [14] S. Dimitrov, “Non-linear distortion cancellation and symbol-based equalization in satellite forward links,” IEEE Trans Wireless Commun, vol. 16, no. 7, pp. 4489–4502, 2017
- [15] R. Thompson, “Nonlinear Distortion Detection using DOCSIS Spectra”, SCTE Conference, Philadelphia, 2016
- [16] C.-L. Wang and Y. Ouyang, “Low-complexity selected mapping schemes for peak-to-average power ratio reduction in OFDM systems,” IEEE Transactions on signal processing, vol. 53, no. 12, pp. 4652–4660, 2005.
- [17] I. Yoffe and D. Wulich, “Predistorter for mimo system with nonlinear power amplifiers,” IEEE Transactions on Communications, vol. 65, no. 8, pp. 3288–3301, 2017.
- [18] J. Kim and K. Konstantinou, “Digital predistortion of wideband signals based on power amplifier model with memory,” Electronics Letters, vol. 37, no. 23, pp. 1417–1418, 2001.
- [19] L. Ding, G. T. Zhou, D. R. Morgan, Z. Ma, J. S. Kenney, J. Kim, and C. R. Giardina, “A robust digital baseband predistorter constructed using memory polynomials,” IEEE Transactions on communications, vol. 52, no. 1, pp. 159–165, 2004.
- [20] M. Rawat, K. Rawat, and F. M. Ghannouchi, “Adaptive digital predistortion of wireless power amplifiers/transmitters using dynamic real valued focused time-delay line neural networks,” IEEE Transactions on Microwave Theory and Techniques, vol. 58, no. 1, pp. 95–104, 2010.
- [21] F. Mkadem and S. Boumaiza, “Physically inspired neural network model for rf power amplifier behavioral modeling and digital predistortion,” IEEE Transactions on Microwave Theory and Techniques, vol. 59, no. 4, pp. 913–923, 2011.
- [22] S. Dimitrov, “Non-linear distortion cancellation and symbol-based equalization in satellite forward links,” IEEE Trans Wireless Communications, vol. 16, no. 7, pp. 4489–4502, 2017.
- [23] D. J. Sebald and J. A. Bucklew, “Support vector machine techniques for nonlinear equalization,” IEEE Transactions on Signal Processing, vol. 48, no. 11, pp. 3217–3226, 2000.

- [24] S. Chen, B. Mulgrew, and P. M. Grant, "A clustering technique for digital communications channel equalization using radial basis function networks," *IEEE Transactions on neural networks*, vol. 4, no. 4, pp. 570–590, 1993.
- [25] B. Li, C. Zhao, M. Sun, H. Zhang, Z. Zhou, and A. Nallanathan, "A Bayesian approach for nonlinear equalization and signal detection in millimeter-wave communications," *IEEE Transactions on Wireless Communications*, vol. 14, no. 7, pp. 3794–3809, 2015.
- [26] M. Rawat, K. Rawat, and F. M. Ghannouchi, "Adaptive digital predistortion of wireless power amplifiers/transmitters using dynamic real valued focused time-delay line neural networks," *IEEE Transactions on Microwave Theory and Techniques*, vol. 58, no. 1, pp. 95–104, 2010.
- [27] PRNewswire, Small Cell Power Amplifier Market - Global Industry Analysis, Size, Share, Growth, Trends and Forecast 2017 – 2025, <https://www.prnewswire.com/news-releases/small-cell-power-amplifier-market---global-industry-analysis-size-share-growth-trends-and-forecast-2017---2025-300546518.html>, October 31, 2017
- [28] B. Li, C. Zhao, M. Sun, H. Zhang, Z. Zhou, and A. Nallanathan, "A Bayesian approach for nonlinear equalization and signal detection in millimeter-wave communications," *IEEE Transactions on Wireless Communications*, vol. 14, no. 7, pp. 3794–3809, 2015.
- [29] M. Rawat, K. Rawat, and F. M. Ghannouchi, "Adaptive digital predistortion of wireless power amplifiers/transmitters using dynamic real valued focused time-delay line neural networks," *IEEE Transactions on Microwave Theory and Techniques*, vol. 58, no. 1, pp. 95–104, 2010.
- [30] D.-C. Park and T.-K. J. Jeong, "Complex-bilinear recurrent neural network for equalization of a digital satellite channel," *IEEE Transactions on Neural Networks*, vol. 13, no. 3, pp. 711–725, 2002.
- [31] A. Uncini, L. Vecchi, P. Campolucci, and F. Piazza, "Complex-valued neural networks with adaptive spline activation function for digital-radiolinks nonlinear equalization," *IEEE Transactions on Signal Processing*, vol. 47, no. 2, pp. 505–514, 1999.
- [32] J. Staudinger, J.-C. Nanan, and J. Wood, "Memory fading Volterra series model for high power infrastructure amplifiers," in *Radio and Wireless Symposium (RWS)*, 2010 IEEE. IEEE, 2010, pp. 184–187.
- [34] A. Zhu, J. C. Pedro, and T. J. Brazil, "Dynamic deviation reduction based Volterra behavioral modeling of rf power amplifiers," *IEEE Transactions on microwave theory and techniques*, vol. 54, no. 12, pp. 4323–4332, 2006.
- [35] L. Guan and A. Zhu, "Simplified dynamic deviation reduction-based Volterra model for Doherty power amplifiers," in *Integrated Nonlinear Microwave and Millimeter-Wave Circuits (INMMIC)*, 2011 Workshop on. IEEE, 2011, pp. 1–4.

# Virtualization and Edge Compute Evolution in Cable

A Technical Paper prepared for SCTE•ISBE by

**Andrii Vladyka**

Technical Product Manager, Cable Access  
Harmonic Inc.  
2590 Orchard Parkway, San Jose, CA 95131  
+1 408 542 2559  
Andrii.Vladyka@harmonicinc.com

**Asaf Matatyaou**

Vice President, Solutions and Product Management, Cable Access  
Harmonic Inc.  
2590 Orchard Parkway, San Jose, CA 95131  
+1 408 542 2559  
Asaf.Matatyaou@harmonicinc.com

# Table of Contents

| Title                                                                    | Page Number |
|--------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                     | 3           |
| 2. Cloud-Native vCMTS Deployment Evolution .....                         | 4           |
| 2.1. Edge Deployment.....                                                | 4           |
| 2.2. CMTS in a CSP Infrastructure.....                                   | 6           |
| 2.3. Hybrid Deployments of Virtualized Cloud-Native Access Platform..... | 7           |
| 2.4. Cloud-Native Virtualized Access Platform and Deep Edge Compute..... | 8           |
| 3. Conclusion.....                                                       | 10          |
| Abbreviations .....                                                      | 10          |
| Bibliography & References.....                                           | 11          |

## List of Figures

| Title                                                                                                       | Page Number |
|-------------------------------------------------------------------------------------------------------------|-------------|
| Figure 1 - Cloud-Native vCMTS Deployment Scenarios .....                                                    | 4           |
| Figure 2 - Disaggregated Virtualized Cloud-Native Access Platform as Stand-Alone Edge Compute Cluster ..... | 5           |
| Figure 3 - Disaggregated Virtualized Cloud-Native Access Platform in Public Cloud .....                     | 6           |
| Figure 4 - Example of Hybrid Deployment of Virtualized Cloud-Native Access Platform.....                    | 8           |
| Figure 5 - Disaggregated Virtualized Access Platform and Deep Edge Compute .....                            | 9           |



# 1. Introduction

Traditionally, the CMTS market was dominated by hardware-based CMTS appliances, with no alternatives available. During the last few years, however, the cable industry took a leap forward and started adopting software-based CMTS solutions<sup>[4]</sup>. Disruptive benefits, long-term potential and deployment scenarios of such software-based CCAP solutions are analyzed in detail in References 1 and 2, while Reference 3 outlines real-world experiences and lessons learned from deploying virtual CMTS in centralized and distributed access architecture (DAA) environments.

This paper analyzes the background and deployment scenarios of a virtualized CMTS (vCMTS) and considers future development of disaggregated vCMTS from the standpoint of leveraging distributed edge compute. We also explore the notion of cluster multitenancy and its potential deployment scenarios in private, public and hybrid cloud.

Here and throughout the paper, we adhere to the following terminology:

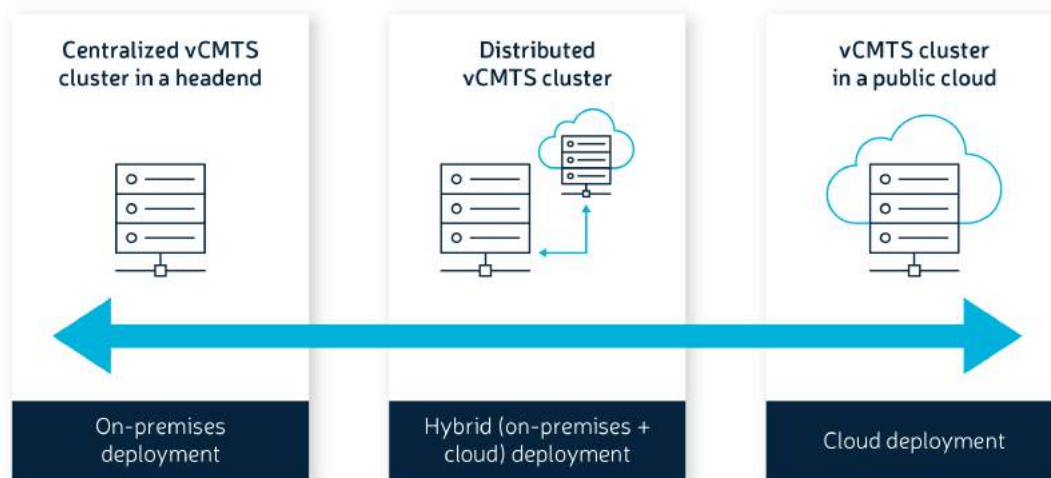
- Cluster – a set of compute nodes that can be viewed as a single system. Cluster nodes are connected with each other via a common converged interconnect network (CIN) and managed by specialized platform management software.
- Pod – a group of one or more software containers, orchestrated by Kubernetes<sup>[16]</sup>.
- Infrastructure (infra) pod – contains applications that perform a management role in the system. The resources and functions provided by infra pods are typically shared by multiple data plane pods. The examples of infrastructure pods are pods running databases used for storing internal system states as well as pods providing interfaces for interaction with users (e.g., command-line interface (CLI), web UI) and toward OSS/BSS applications (e.g., via SNMP, IPDR and other protocols).
- vCMTS pod – a collection of containers implementing DOCSIS<sup>[14]</sup> processing (excluding the PHY layer).
- Data plane pod – a more generic term for the vCMTS pod. A data plane pod contains software components performing data plane traffic processing for an access technology of choice (e.g., DOCSIS, PON, wireless). It should be noted that PHY layer processing is implemented in purpose-built hardware appliances (e.g., DOCSIS 3.1 Remote PHY Device) tailored for communication with the user equipment (e.g., DOCSIS 3.1 cable modem) over specific media type (e.g., coax cable or optical fiber).
- Multitenant cluster – sharing compute resources for infrastructure pods and data plane pods serving different types of applications (e.g., DOCSIS, PON, wireless), while providing required level of resources and components isolation.
- Hybrid cluster deployment – combining local compute resources and compute resources from cloud service provider (CSP) in the same logical cluster.
- Virtualized cloud-native access platform – software platform that follows the cloud-native paradigm<sup>[15]</sup> and provides facilities for life cycle management of compute resources and software applications running on top of them. This type of platform is agnostic to the type of applications (e.g., data plane pods) deployed on top of it.

In the subsequent sections, we look in more detail at a range of options that become available to operators when migrating to a cloud-native vCMTS.

## 2. Cloud-Native vCMTS Deployment Evolution

CMTS is not the only product that transformed dramatically over the last few years. Studies<sup>[11]</sup> show that both public and private cloud adoption grows, with larger enterprises increasing their focus on public cloud. Moving services to the public cloud entails an operations paradigm shift. It is possible to incorporate software-based vCMTS deployment and operations into cloud-centric business processes and automation tools. If done properly, it allows operators to take service agility to the next level, dramatically reducing time to market for rolling out new services and transforming the CMTS from a purpose-built and rigid resource in a headend to an on-demand scalable service.

Due to its disaggregated architecture, cloud-native vCMTS provides full flexibility when selecting deployment type – from on premises to public cloud (Figure 1 - Cloud-Native vCMTS Deployment Scenarios).



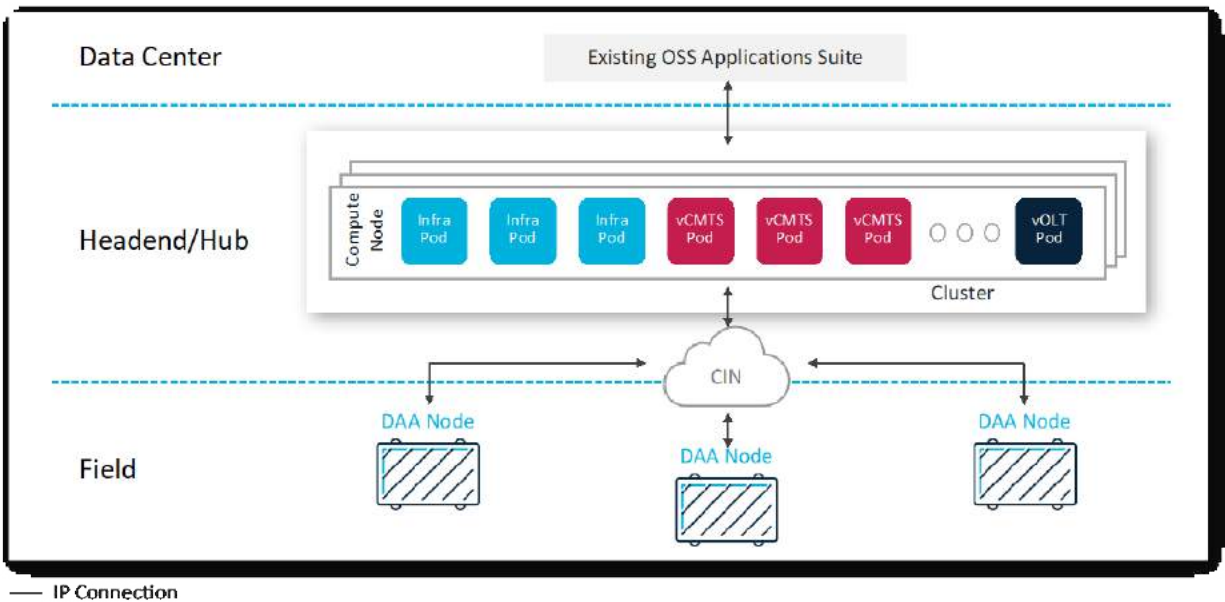
**Figure 1 - Cloud-Native vCMTS Deployment Scenarios**

Another dimension to look at is types of access technologies that can be deployed on the same cluster and managed in a uniform way: i.e., cluster compute resources that can be shared between different workload types (e.g., vCMTS data plane pods and virtual PON (vPON) data plane pods). It introduces the notion of cluster multitenancy, where the same virtualized cloud-native access platform accommodates data plane pods for different types of access technologies. We will use the more generic “access platform” term throughout the paper when it is required to refer to the workload-agnostic nature of the deployment type.

In the following sections, we look into different scenarios of disaggregated vCMTS deployment, from on-premises to public cloud.

### 2.1. Edge Deployment

One side of the spectrum of options when deploying cloud-native vCMTS is running it on dedicated compute and management cluster nodes located in the headends and remote hubs, the same way it is done for legacy hardware-based CMTS appliances.



**Figure 2 - Disaggregated Virtualized Cloud-Native Access Platform as Stand-Alone Edge Compute Cluster**

In this scenario, individual clusters work independently of each other, and while cluster scale (e.g., number of compute nodes) can vary from site to site, every cluster acts as a single deployment unit. All resources required for CMTS functionality (e.g., data plane workloads, generic control plane (GCP) Core, interfaces to OSS applications) are unique per cluster. While at first sight it looks like yet another CMTS, there are clear benefits of migrating to a cloud-native architecture, even in a stand-alone edge deployment:

1. Increase reliability by leveraging cloud-native best practices and tools for high availability (HA).
2. Reduce operating domain to one or several service groups by isolating vCMTS resources with Docker containers.
3. Solve scaling issues by simple addition of compute nodes, with growth in bandwidth as well as the number of subscribers and/or service groups.
4. Significantly reduce power and cooling requirements, space footprint and wiring complexity in the headend.
5. Ability to share cluster compute resources between different types of workloads (data plane pods). As such, the same cluster can accommodate, for example, vCMTS and vPON pods, while new emerging types of data plane pods might be added to a cluster in the future.

Horizontal scaling in stand-alone edge deployment is achieved by scaling the number of vCMTS pods (or other data plane pods) on existing compute nodes, as well as by adding more compute nodes proportionally to the number of connected devices. The contributing factors into cluster scaling are:

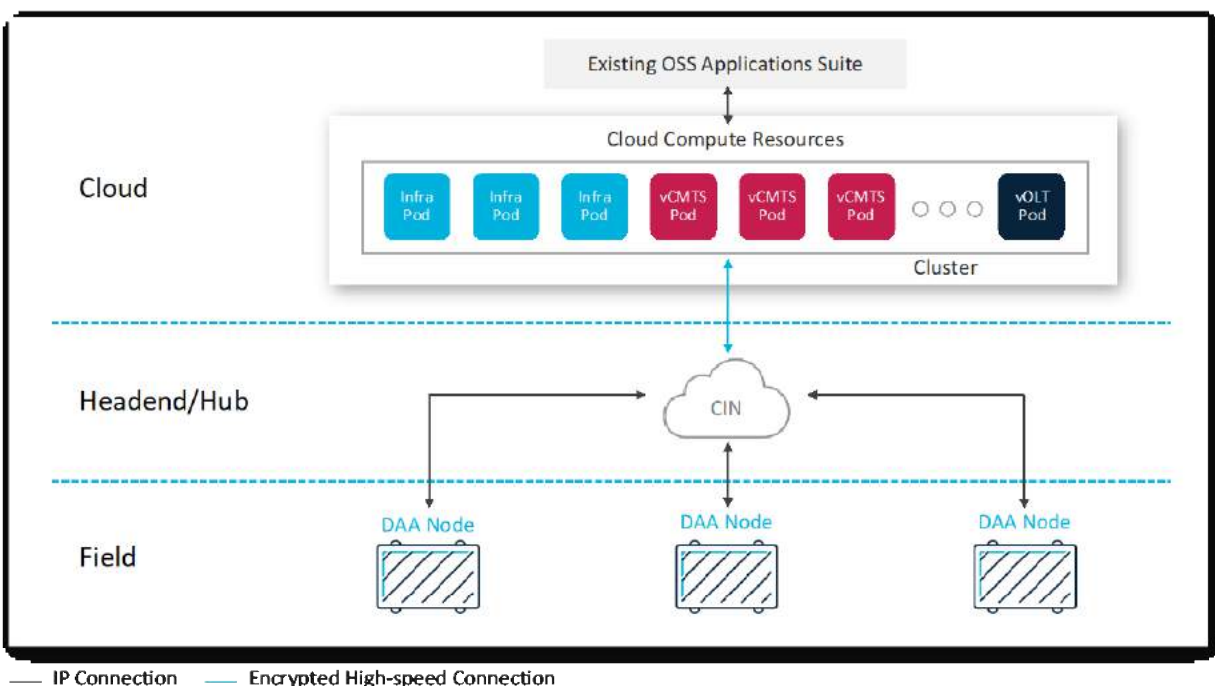
1. The number of connected Remote PHY Devices (RPDs).
2. The number of connected Cable Modems (CMs) and Customer Premises Equipment (CPE) units.
3. Peak traffic load.

Cluster scaling implies increasing the number of data plane pods, while the number of infra pods remains constant. Keeping the amount of infra resources constant with the growing number of data plane pods further improves the efficiency of a virtualized cloud-native access platform.

Such a deployment scenario is most suitable for cable operators who would like to follow the traditional way of operating CMTSs and do not plan to take steps toward building network function virtualization (NFV) infrastructure and/or network automation or run some of the vCMTS components in the hybrid cloud.

## 2.2. CMTS in a CSP Infrastructure

Another potential scenario of cloud-native vCMTS deployment is to completely decouple it from hardware resources and run it in a public cloud on compute resources provided by a CSP.



**Figure 3 - Disaggregated Virtualized Cloud-Native Access Platform in Public Cloud**

This type of deployment implies that all control plane and data plane vCMTS components run on compute resources in the public cloud. Remote PHY protocol (e.g., downstream external PHY interface (DEPI), upstream external PHY interface (UEPI)) tunnels are established between the vCMTS data plane workloads and R-PHY nodes, and all data processing is performed in the public cloud. The operations support systems (OSS) applications suite may remain in the operator's data center, or it can be deployed in the public cloud as well. The factors that contribute to the decision-making process are:

1. Network total cost of ownership (TCO) optimization.
2. Migration path for legacy OSS applications.
3. Regulations compliance.

Such a solution requires reliable and secure connections between the R-PHY nodes and vCMTS components running in the cloud. It is achieved in the following ways:

1. Using facilities provided by CSPs, such as encryption gateways that enable secure high-speed data transfer between the operator's data center and public cloud infrastructure<sup>[5]</sup>.
2. By leveraging a suite of security features defined for RPD and CCAP Core Mutual Authentication<sup>[6]</sup> and DOCSIS Baseline Privacy<sup>[7]</sup>.
3. Additional security measures might be applicable to other types of access technologies deployed on a multitenant cluster.

Moving all data and control plane resources to a public cloud can be considered by operators pursuing the following goals:

1. Minimize power consumption and space footprint in operator's facilities. In this scenario, only routing/switching equipment remains in the headend/hub.
2. Accelerate and streamline service deployment and nodes splits. vCMTS resources can scale on demand in a matter of minutes/hours in response to business needs.
3. Move to advanced high availability (HA) scenarios with geo-redundancy, with different components of disaggregated vCMTS deployed in different locations. For example, it should be possible to deploy critical system components in geographically distributed availability zones in public cloud and build redundant network paths between hub/headend and CSP infrastructure.
4. Unify vCMTS resource provisioning with existing business processes, such as service provisioning, and leverage cloud-native architecture for achieving CMTS-as-a-service deployment scenarios.

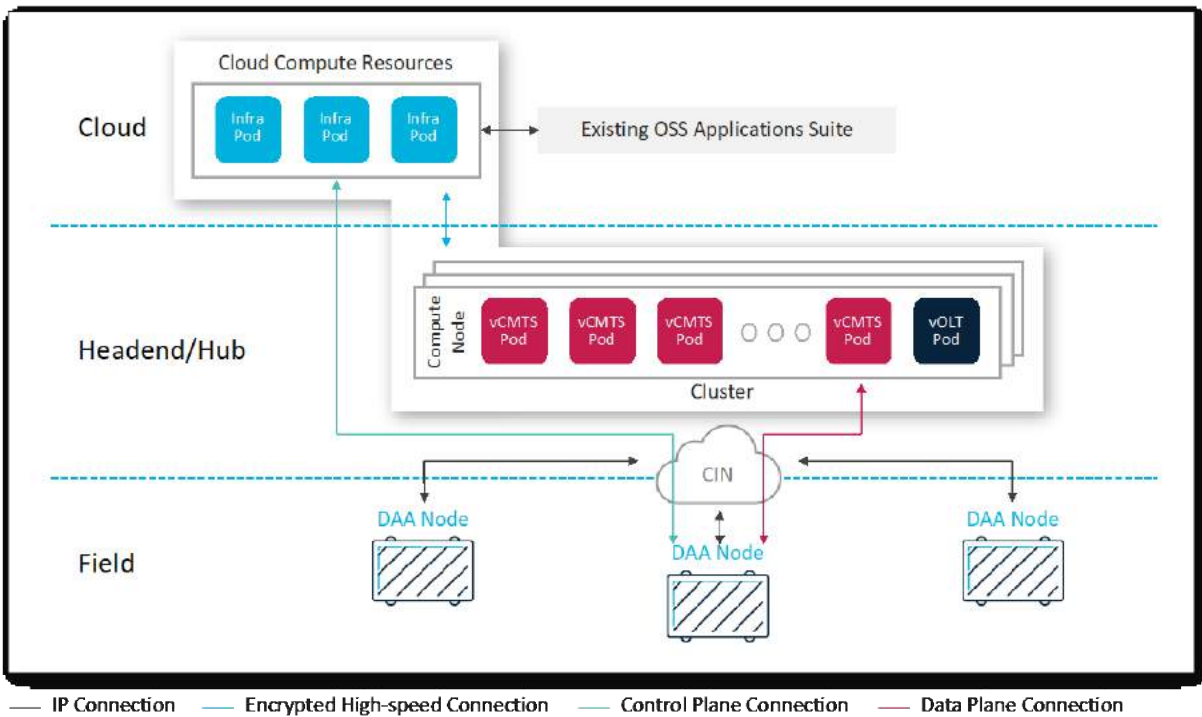
The CMTS scaling factors described for the on-premises edge deployments are also valid for a public cloud deployment scenario. Additional benefits in the latter case include:

1. Improved flexibility for horizontal and vertical node scaling.
2. Reduced time to market when adding more compute nodes to the cluster.
3. Transfer of compute resources life cycle management and associated cost from the operator's domain of responsibility to a CSP.

It should be noted that distance and latency between provider edge and cloud-based facilities should be considered when planning vCMTS deployment scenarios in the public cloud. Another important point is that measures to achieve advanced HA capabilities are not limited to protecting components deployed in public cloud and should also include necessary steps to protect system components in headend/hub and in the field (e.g., CIN redundancy, RPD network connectivity redundancy).

### **2.3. Hybrid Deployments of Virtualized Cloud-Native Access Platform**

In a broad sense, hybrid cloud platforms from almost all the major infrastructure vendors not only manage clusters running on premises and in their own cloud platforms, but any cluster including those that are deployed in other cloud environments<sup>[8]</sup>. Due to its cloud-native packaging, disaggregated vCMTS provides enough flexibility to deploy it in hybrid scenarios, where data plane pods run on local edge compute nodes, while infra pods are deployed in private or public cloud. It allows operators to fill the gap between on-premises deployment and cloud-based deployment with other options that combine the benefits of both.



**Figure 4 - Example of Hybrid Deployment of Virtualized Cloud-Native Access Platform**

Hybrid deployment of the disaggregated virtualized cloud-native access platform is suitable for use cases wherein data plane processing is kept in provider edge facilities, orchestrated by virtualized control plane resources from public or private cloud. The benefits of such deployment type include but are not limited to:

1. Minimized latency between end-user equipment and data plane pods.
2. Possibility to deploy infrastructure and data plane resources in existing k8s clusters, with unified orchestration of vCMTS workloads and third-party workloads.
3. Ability to handle high-availability scenarios, possibly with geo-redundancy provided by a CSP, for control plane components and data plane components in a uniform way. From a pod orchestration standpoint, cloud compute resources and on-premises compute resources are considered to be part of the same cluster.

From a horizontal scaling point of view, the hybrid deployment scenario combines the benefits of on-premises and public cloud deployment scenarios.

There are, of course, other options for deploying disaggregated vCMTS in an operator's edge facilities as well as in a public and private cloud environment. In the following section, we'll look at how the concept of cloud-native disaggregated access platform allows operators to leverage emerging deployment scenarios.

## 2.4. Cloud-Native Virtualized Access Platform and Deep Edge Compute

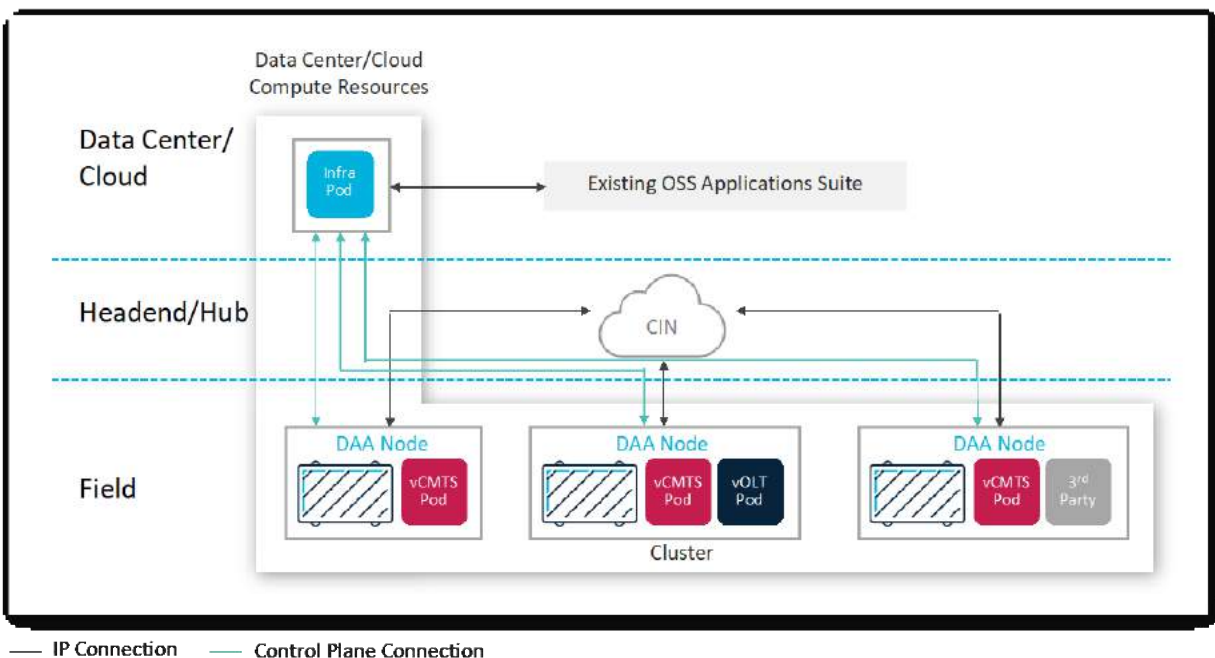
So far, we discussed centralized compute resources available to cable operators in a public and private cloud, as well as edge compute resources distributed across headends and hubs. As equipment deployed on the outside plant becomes more intelligent and provides more processing power, we can leverage its

capabilities and extend the range of the compute resources available to cable operators to so-called “deep edge” compute. Therefore, deployment concepts we discussed so far can be considered as a superset of CableLabs’ Flexible MAC architecture (FMA)<sup>[9]</sup> that specifies interoperable management, control and data plane interfaces on the digital fiber between the core components and a series of devices, regardless of compute locality.

A virtualized cloud-native access platform allows operators to control the following aspects of FMA deployments:

1. Locality of certain vCMTS components, based on the distance/latency requirements and constraints. Data plane pods can be placed on the compute resources located in the outside plant (so-called “deep edge” deployment), while infrastructure components can be deployed in a centralized fashion.
2. Flexible separation of individual vCMTS components (microservices) and their subsequent placement on the compute resources in the data center, headend/hub or outside plant.
3. Backward compatibility with existing OSS applications suite.
4. Fast time to market and ability to add new services and applications on top of the cluster that consists of general-purpose compute resources.

When deployed as part of a cloud-native virtualized access platform, a group of FMA remote MAC core units (RMC) and/or virtual core (vCore) units in effect becomes a pool of general-purpose compute nodes orchestrated by Kubernetes.



**Figure 5 - Disaggregated Virtualized Access Platform and Deep Edge Compute**

RMD/RMC/vCore can be considered as a generic compute resource that can be utilized, in particular, for DOCSIS MAC layer processing combined with purpose-built hardware that implements, for example, the DOCSIS PHY layer. In such a case, it is possible to push vCMTS pods to the very network edge for vCMTS software-based packet processing, while leveraging best industry practices for cloud-native

software deployment and life-cycle management. It also opens the door for utilizing RMD compute resources for other types of workloads, such as:

1. IoT edge processing.
2. Sharing common generic compute resources in RMD for different types of data plane pods (e.g., perform DOCSIS and wireless data plane processing on the same compute resources in RMD).
3. Hosting applications that can benefit from deep edge deployment: caching servers, speed test applications, VR-specific applications, machine-to-machine communication applications.

Horizontal scaling and infra resource-sharing concepts described earlier are also valid for multitenant clusters. Potential benefits from placing DOCSIS and third-party workloads on deep edge compute nodes include:

1. Reduced end-to-end network latency, which is inherited from the fact that compute resources are located in physical proximity to the customer end devices.
2. Maximized resource utilization that emerges from using generic edge compute resources for DOCSIS and non-DOCSIS applications.
3. Network traffic reduction and savings on bandwidth between deep edge devices and the central cloud.

### 3. Conclusion

Compute resources available for cable operators span from centralized compute in public and private cloud to the edge compute in headend/hub and deep edge compute in an outside plant. Cloud-native virtualization enables operators to leverage distributed compute resources in a uniform way.

A virtualized cloud-native and disaggregated vCMTS provides a wide range of options for deployment:

1. Traditional on-premises edge deployments in headend/hub.
2. Cloud-based deployments in private cloud and in CSP infrastructure.
3. Hybrid deployments, wherein certain vCMTS components remain in operator's facilities, while other components get pushed to a CSP infrastructure.
4. Deep edge deployments, where some of the cluster compute resources are located in the outside plant.

In all deployment scenarios, cluster resource and application life cycle management are performed in a uniform way:

1. Compute resources are considered to be general-purpose.
2. Heterogeneous workloads (different types of infrastructure, data plane and third-party pods) are placed on general-purpose compute resources, taking into account application locality preferences (e.g., cloud/data center, network edge or deep network edge).
3. Horizontal scaling is achieved by increasing the number of workloads (data plane pods) and by adding more compute resources (in the cloud or on premises).

The combination of the available compute resources and the flexibility of cloud-native access platform allows operators to deploy different types of access networks in a uniform and operations efficient way. Implementing as many network functions as possible in software and minimizing the number of purpose-built hardware appliances on the network provides a bridge to the future that was never more unknown than now.

## Abbreviations

|      |                                 |
|------|---------------------------------|
| CCAP | converged cable access platform |
|------|---------------------------------|



|        |                                                 |
|--------|-------------------------------------------------|
| CIN    | converged interconnect network                  |
| CLI    | command-line interface                          |
| CM     | cable modem                                     |
| CMTS   | cable modem termination system                  |
| CPE    | customer premises equipment                     |
| CSP    | cloud service provider                          |
| DAA    | distributed access architecture                 |
| DEPI   | downstream external PHY interface               |
| DOCSIS | Data Over Cable Service Interface Specification |
| FMA    | Flexible MAC Architecture                       |
| GCP    | generic control plane                           |
| HA     | high availability                               |
| HFC    | hybrid fiber-coax                               |
| IoT    | Internet of things                              |
| IPDR   | Internet Protocol Detail Record                 |
| MAC    | media access control                            |
| NFV    | network function virtualization                 |
| OLT    | optical line terminal                           |
| OSS    | operations support systems                      |
| PHY    | physical layer                                  |
| PON    | passive optical network                         |
| RMC    | remote MAC core                                 |
| RPD    | remote PHY device                               |
| SCTE   | Society of Cable Telecommunications Engineers   |
| SNMP   | Simple Network Management Protocol              |
| TCO    | total cost of ownership                         |
| UEPI   | upstream external PHY interface                 |
| UI     | user interface                                  |
| vCMTS  | virtual CMTS                                    |
| vCore  | virtual Core                                    |
| vPON   | virtual PON                                     |

## Bibliography & References

- [1] Matatyaou, Asaf. Transforming the HFC Access Network with a Software-Based CCAP. Publication. San Jose: Harmonic, 2015. Web.
- [2] Matatyaou, Asaf. Real-World Deployment of a Virtual Cable Hub. Publication. San Jose: Harmonic, 2017. Web.
- [3] Matatyaou, Asaf. Practical Lessons of a DAA Deployment with a Virtualized CMTS. Publication. San Jose: Harmonic, 2017. Web.
- [4] Robuck, Mike. Comcast executives on lessons learned from deploying vCMTS. October 29, 2018. Web. <https://www.fiercetelecom.com/telecom/comcast-execs-lessons-learned-from-deploying-vcmts>
- [5] Barr, Jeff. New – Use an AWS Transit Gateway to Simplify Your Network Architecture. November 26, 2018. Web. <https://aws.amazon.com/blogs/aws/new-use-an-aws-transit-gateway-to-simplify-your-network-architecture/>
- [6] Remote PHY Specification, CM-SP-R-PHY-I14-200323, March 3, 2020, Cable Television Laboratories, Inc.

- [7] Baseline Privacy Plus Interface Specification, CM-SP-BPI+-C01-081104, March 11, 2008, Cable Television Laboratories, Inc.
- [8] MSV, Janakiram. How Kubernetes Has Changed The Face Of Hybrid Cloud. December 16, 2019. Web. <https://www.forbes.com/sites/janakirammsv/2019/12/16/how-kubernetes-has-changed-the-face-of-hybrid-cloud>
- [9] Flexible MAC Architecture (FMA) System Specification, CM-SP-FMA-SYS-D04-200218, February 18, 2020, Cable Television Laboratories, Inc.
- [10] Milinkovich, Mike. K8s at the Edge – Some Context on the New Kubernetes IoT Working Group. September 26, 2018. Web. <https://eclipse-foundation.blog/2018/09/26/k8s-at-the-edge/>
- [11] Weins, Kim. Cloud Computing Trends: 2018 State of the Cloud Survey. February 13, 2018. Web. <https://www.flexera.com/blog/cloud/2018/02/cloud-computing-trends-2018-state-of-the-cloud-survey/>
- [12] Computer cluster. <https://en.wikipedia.org/>, April 15, 2020. Web. [https://en.wikipedia.org/wiki/Computer\\_cluster](https://en.wikipedia.org/wiki/Computer_cluster)
- [13] What Is NFV Infrastructure (NFVI)? Definition. <https://www.sdxcentral.com/>, April 27, 2016. Web. <https://www.sdxcentral.com/networking/nfv/definitions/nfv-infrastructure-nfvi-definition/>
- [14] MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I20-200407, April 6, 2020, Cable Television Laboratories, Inc.
- [15] CNCF Cloud Native Definition v1.0. Web. <https://github.com/cncf/toc/blob/master/DEFINITION.md>
- [16] Kubernetes. Web. <https://github.com/kubernetes/kubernetes>

# **Path To Gigabit Fixed Wireless Access**

## **A Review Of Fixed Wireless Access Technology And Economics**

A Technical Paper prepared for SCTE•ISBE by

### **Venkata Somi Alapati**

VP Business Development and Wireless Solutions  
Ericsson  
6300 Legacy Drive Plano, TX 75024  
(972) 955-3718  
somi.alapati@ericsson.com

### **Kashif Shakil**

Customer Solutions Sales Director  
Ericsson  
6300 Legacy Drive Plano, TX 75024  
(972) 679-3737  
Kashif.shakil@ericsson.com

### **Craig Schwechel**

Principal Consultant  
Incode Consulting  
6300 Legacy Drive Plano, TX 75024  
(434) 242-2843  
craig.schwechel@incodeconsulting.com

# Table of Contents

| <b>Title</b>                                     | <b>Page Number</b> |
|--------------------------------------------------|--------------------|
| Table of Contents .....                          | 2                  |
| 1. Introduction.....                             | 4                  |
| 2. Reasons for FWA Momentum.....                 | 5                  |
| 3. Fixed Wireless Access Network.....            | 6                  |
| 4. Spectrum Options for FWA .....                | 7                  |
| 5. Starting Up and Evolving FWA Network .....    | 8                  |
| 6. 3GPP Data Rate Evolution.....                 | 9                  |
| 7. Beamforming Technology .....                  | 10                 |
| 8. CBRS for FWA .....                            | 11                 |
| 9. Massive MIMO for FWA .....                    | 12                 |
| 10. FWA with 5G .....                            | 15                 |
| 11. Improving 5G mmWave Reach .....              | 16                 |
| 12. Summary of Technical Section .....           | 20                 |
| 13. FWA Business Considerations.....             | 21                 |
| 14. FWA economics .....                          | 21                 |
| 15. CBRS Business Case Example .....             | 23                 |
| 16. mmWave and CBRS Business Case Example.....   | 23                 |
| 17. FWA business models .....                    | 24                 |
| 18. CBRS FWA Field Experiences.....              | 24                 |
| 19. mmWave and CBRS Field Trial experiences..... | 27                 |
| Conclusions.....                                 | 27                 |
| Abbreviations .....                              | 28                 |
| Bibliography & References.....                   | 28                 |

## List of Figures

| <b>Title</b>                                                                     | <b>Page Number</b> |
|----------------------------------------------------------------------------------|--------------------|
| Figure 1: Fixed Wireless Access network .....                                    | 6                  |
| Figure 2: LTE throughput ladder .....                                            | 10                 |
| Figure 3: Digital Beamforming .....                                              | 11                 |
| Figure 4: Analog Beamforming .....                                               | 11                 |
| Figure 5: Coverage under broadcast beam .....                                    | 12                 |
| Figure 6: Beamforming with SU-MIMO .....                                         | 13                 |
| Figure 7: Beamforming with MU-MIMO .....                                         | 14                 |
| Figure 8: Comparison of 4T4R and 64T64R user throughput and coverage .....       | 15                 |
| Figure 9: Combining mmWave and mid band.....                                     | 17                 |
| Figure 10: Dual Connectivity and Carrier Aggregation .....                       | 18                 |
| Figure 11: Dual Connectivity Data Bearer Architecture .....                      | 18                 |
| Figure 12: mmWave and CBRS LOS Bandwidth vs Distance.....                        | 20                 |
| Figure 13: mmWave and CBRS NLOS Bandwidth vs Distance .....                      | 20                 |
| Figure 14: Downlink throughput under different RF conditions.....                | 25                 |
| Figure 15: Setting for multiple user MIMO testing in ideal radio conditions..... | 25                 |

|                                                                 |    |
|-----------------------------------------------------------------|----|
| Figure 16: FWA with beamforming and MU-MIMO in rural area ..... | 27 |
|-----------------------------------------------------------------|----|

## List of Tables

| <b>Title</b>                                                                                | <b>Page Number</b> |
|---------------------------------------------------------------------------------------------|--------------------|
| Table 1: Broadband access technology comparison .....                                       | 4                  |
| Table 2: Spectrum for FWA.....                                                              | 7                  |
| Table 3: Results of throughput testing with 16 layer MU-MIMO in good radio conditions ..... | 26                 |

# 1. Introduction

Consumer demand for home broadband access continues to be strong. Network operators are looking to grow their revenue by expanding services they offer. Home broadband offers that revenue growth opportunity.

1. Operators can expand home broadband services to unserved and underserved communities. Rural markets are one example of underserved communities.
2. For some network operators, even communities served by existing Internet service providers are attractive target for business expansion. These operators bring something new to the market, like higher data rate or lower cost or some other valuable feature.

When thinking about greenfield opportunities, operators must consider technology options available to them. Below we compare technology options for providing home broadband service.

**Table 1: Broadband access technology comparison**

| Technology | Unserved/Under-served (mostly rural markets)                                                                                                                                                  | Urban and suburban markets                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DSL        | It is expensive to build out DSL plant in rural areas since there is limitations on how far from central office can the plant be extended. Laying new copper is also an expensive proposition | DSL is less competitive with DOCSIS or fiber in more urban markets. For instance, maximum rates offered by VDSL2 could be greater than 100 Mbps with a range of around 500m from DSLAM node.<br><br>G.fast promises higher data rates for shorter straight loops. For instance, 600 Mbps for 200m distance. Speed limitation start kicking for longer distances. Also copper is hard to maintain.<br><br>There is no viable path for bit rate evolution beyond that. |
| DOCSIS HFC | May not be cost effective to expand cable plant into rural areas                                                                                                                              | DOCSIS 4.0 FDX or ESD offers 10 Gbps shared downstream capacity. It is an attractive option but fiber and 5G offerings could disrupt DOCSIS.                                                                                                                                                                                                                                                                                                                         |
| Fiber      | Running new fiber to rural communities can get very expensive                                                                                                                                 | Fiber is the leading medium for data transmission with virtually unlimited bandwidth. However, green field FTTH installations may not be economically viable. Similarly upgrading HFC or DSL plants to FTTH may only be feasible for selected communities                                                                                                                                                                                                            |

|                        |                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3GPP wireless (LTE/5G) | Since there is no need to run copper or fiber over long distances, 3GPP wireless constitutes an economic option to provide rural broadband service as long as the capacity needs and path loss are acceptable.      | <p>New technologies like massive MIMO and availability of more spectrum in 5G mmWave bands enhance available capacity and range of a single cell site. This improves business feasibility of fixed wireless access FWA services.</p> <p>New operators can disrupt existing DOCSIS and fiber-based offerings.</p> <p>Established operators should study leveraging 3GPP wireless to defend and grow their business.</p> |
| Proprietary wireless   | Proprietary wireless technologies are by definition not standards based, they may be limited to one vendor, have smaller ecosystem, an uncertain roadmap & lack of economies of scale and limited long term support | Whether rural or urban, proprietary technologies face similar hurdles.                                                                                                                                                                                                                                                                                                                                                 |

In this paper, we focus on 3GPP based FWA, as a promising technology for future home broadband access. We investigate the following in the coming pages:

1. Can today's wireless technology support FWA?
2. Is there a viable business case for FWA as compared to alternatives (DOCSIS, FTTH)?
3. Do any network operators have profitable FWA home broadband business?

## 2. Reasons for FWA Momentum


Several factors are coming together, boosting industry push for FWA:

1. Technology advances: Higher order modulation, massive MIMO, beamforming, carrier aggregation, multiple user MIMO are some of the features in LTE that are improving spectral efficiency of LTE. This has the dual effect of improving system capacity, but also cell edge bit rates – suitable for FWA. 5G adds further improvements to spectral efficiency by leaner air interface design.
2. Spectrum is the lifeline of wireless access. New spectrum has become available in mid band and higher bands. This spectrum offers wider channels and thus order of magnitude higher throughput than traditional cellular and PCS bands. LTE also allows use of unlicensed or semi-licensed frequencies in 5 GHz band and 3.5 GHz CBRS in conjunction with licensed bands to further uplift data rates.
3. As more and more network operators include FWA in their consumer offerings, a device and terminal ecosystem has developed. FWA offerings started in developing countries where high speed wireline infrastructure is lacking. That has helped to nurture the FWA ecosystem for everyone. One of the most important factors is formation of global ecosystem driving economies of scale. CPE price is a big factor in overall business case and hence it is important to have global economies of scale.

4. Governments realize the productivity gain and development effects of broadband connection availability to all citizens. Governments at many levels (federal, state, city) want to encourage deployment of high-speed home broadband in their jurisdictions. US federal government CAFII initiative is an example of subsidizing build out of rural broadband networks. CAFII program encourages network operators to ramp up broadband deployments in un/under served, mostly rural or exurban areas. RDOF (Rural Digital Opportunity Fund) focuses on areas unserved at 25/3 in traditional price cap territories and provides \$16B funds to serve ~6M potential home locations with speed tiers of 25/3, 50/5, 100/20, 1000/500 Mbps (Downlink/Uplink).

### 3. Fixed Wireless Access Network

FWA largely reuses MBB network architecture, with similar nodes as mobile broadband sans network features related to terminal mobility.



**Figure 1: Fixed Wireless Access network**

LTE or 5G or dual LTE/5G CPE is installed on the rooftop of the home. This is typically a professional installation. Some operators offer indoor CPEs which are self-installed by the homeowner. Outdoor CPE has high gain antenna and LTE/5G modem to provide connection towards nearest LTE or 5G base station. Operators may provision an FWA GW inside home or integrate it with the outdoor CPE. FWA GW provides home broadband management functionality to the operator.

FWA network comprises of standard 3GPP architecture with an LTE/5G base station and core network. Operator can choose their own deployment strategies. Some considerations below:

1. Dedicated base station and frequencies for FWA or sharing base station and frequencies between FWA, MBB, IoT and other operator services
2. Dedicated core network for FWA or shared with other operator services
3. Physical or virtualized core and radio base station
4. Placement locations of radio and baseband processing of base station, as well as user and control plane components of core network
5. Type of transport between remote (base station) site and aggregation and data center sites.

FWA may require different architectural optimizations different from MBB. For instance, both user plane and control plane design of core network could be different for FWA subscribers. FWA service has fewer users per cell, thus lower control plane load and light-



weight control plane nodes. On the other hand, data consumption is higher in the user plane. A more distributed user plane would be beneficial to improve network performance and reduce transmission costs. Similarly, centralizing baseband processing would have lower utility with reduced pooling gains and might not be worth the increased fronthaul transmission costs.

## 4. Spectrum Options for FWA

Spectrum is the lifeblood of wireless communications. Not all spectrum is the same. Table below summarizes spectrum options commonly used for FWA.

**Table 2: Spectrum for FWA**

| Frequency                    | Benefits                                                                                                                                                                                                                              | Challenges                                                                                                                                                                  | Availability                                                                                                              |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| Mid band FDD (PCS, AWS etc.) | <ul style="list-style-type: none"> <li>– Widely available ecosystem</li> <li>– Licensed</li> <li>– Good propagation</li> </ul>                                                                                                        | <ul style="list-style-type: none"> <li>– Fully utilized in urban and suburban areas</li> <li>– Narrower channel bandwidths</li> </ul>                                       | <ul style="list-style-type: none"> <li>– Now</li> </ul>                                                                   |
| 2.5 GHz TDD                  | <ul style="list-style-type: none"> <li>– ~200 MHz of licensed spectrum</li> <li>– Best propagation amongst TDD spectrum</li> <li>– US ecosystem available today</li> <li>– Highest predictability due to licensed spectrum</li> </ul> | <ul style="list-style-type: none"> <li>– Cost of acquiring spectrum</li> <li>– Majority of spectrum in populated areas is owned</li> </ul>                                  | <ul style="list-style-type: none"> <li>– Now</li> </ul>                                                                   |
| 3.5 GHz CBRS                 | <ul style="list-style-type: none"> <li>– 150 MHz of spectrum</li> <li>– Global LTE ecosystem</li> <li>– Good balance between propagation, power and reuse</li> <li>– Interference managed via SAS (Spectrum Access System)</li> </ul> | <ul style="list-style-type: none"> <li>– Spectrum demand in urban and suburban areas</li> </ul>                                                                             | <ul style="list-style-type: none"> <li>– Late-2019 for GAA and PAL spectrum available by 4Q 2020</li> </ul>               |
| 3.7 GHz C-Band               | <ul style="list-style-type: none"> <li>– 280 MHz of spectrum</li> <li>– Global LTE ecosystem</li> <li>– Better propagation and power compared to CBRS</li> </ul>                                                                      | <ul style="list-style-type: none"> <li>– Cost of acquiring</li> </ul>                                                                                                       | <ul style="list-style-type: none"> <li>– Auction Dec 2020; Deployment 1st phase: Dec 2021, 2nd phase: Dec 2023</li> </ul> |
| 5 GHz                        | <ul style="list-style-type: none"> <li>– 555 MHz of spectrum</li> </ul>                                                                                                                                                               | <ul style="list-style-type: none"> <li>– Propagation challenges - Maximum of 36 dBm EIRP.</li> <li>– Prone to interference due to contention-based access method</li> </ul> | <ul style="list-style-type: none"> <li>– Today</li> </ul>                                                                 |

|                       |                                                                    |                                                                   |                                         |
|-----------------------|--------------------------------------------------------------------|-------------------------------------------------------------------|-----------------------------------------|
| 5.9 GHz-<br>6.425 GHz | - 500 MHz of spectrum                                              | - Following unlicensed framework established for 5 GHz            | - Estimated 2022                        |
| 24 GHz –<br>39 GHz    | - Channels of 100 MHz possible<br>- Carrier agg of 400 MHz or more | - Significant challenges - propagation                            | - Starting mid-2019                     |
| 57 GHz –<br>71 GHz    | - 14 GHz of spectrum<br>- Suitable for point-to-point              | - Significant challenges - propagation and atmospheric absorption | - Partially Today;<br>Partially in 2022 |

Operators providing MBB services could use any of the spectrum above for FWA depending on utilization of their spectrum resources. Greenfield operators could start off with CBRs spectrum in combination with 5 GHz band and perhaps also upcoming mmWave high bands.

Carrier aggregation of LTE/5G bands becomes a critical feature for FWA services:

1. There may not be enough spectrum in one band
2. Spectrum may not be contiguous
3. Combining high bands with lower bands improves cell edge rates and extends the coverage
4. Higher peak and average rates can be offered
5. Higher UL Data Rates may require carrier aggregation with other frequency bands

## 5. Starting Up and Evolving FWA Network

As network traffic grows and as operator offer higher rate broadband services, there will be a need to enhance and upgrade the FWA network. Operator with existing cellular assets could follow a network strategy as below.

1. Use existing MBB cell-site infrastructure for FWA sites, by just adding carriers or slices for FWA. This enables cost effective start up for FWA services
2. Install outdoor CPEs on rooftops of homes or on sides of buildings. Outdoor CPE enable better coverage, longer range and higher system capacity
3. For many existing operators, legacy bands are fully occupied by MBB DL traffic. There may still be capacity left over, specifically in the UL portion of FDD bands that could be given to FWA traffic
4. Operators could upgrade legacy FDD bands to 4T/4R configurations or even FDD massive MIMO to squeeze more network capacity out of legacy FDD bands

5. Operators could add TDD band in 2.5 or 3.5 GHz (with massive MIMO radios). TDD bands are well suited to FWA. Spectrum costs tend to be lower and TDD profiles could match asymmetric downlink heavy nature of home broadband traffic
6. Operators can add mmWave 5G sites in high traffic demand areas to offload FWA traffic from larger macro type sites. mmWave 5G could also be used to provide very high throughput (~ 1 Gbps) service to selected neighborhoods
7. To evolve networks further and to take advantage of 5G's better spectral efficiency and operational ease, operators could upgrade 4G bands to 5G. Since operators need to support both 4G and 5G terminals, there will be a need to operate 5G network together with 4G using schemes like real time spectrum sharing. Real time spectrum sharing allocates spectrum proportionally to 5G and 4G users, as per real time usage demand, without the need to partition spectrum statically and reducing data rates for legacy 4G users
8. Operators could introduce virtualized RAN running on COTS hardware and centralize deployment and management of pieces of FWA RAN functionality especially in Sub-urban and Urban areas.

Greenfield FWA operators that do not have existing spectrum assets could start with step 5 above. Greenfield operators would need to consider their migration strategies to 5G as well. 5G migration that could be accomplished with software upgrade to 4G eNBs and 4G CPEs provide a more compelling option. Even greenfield operators must consider implications of providing 5G and 4G services on the same spectrum, as their user base migrates over time from 4G to 5G.

## **6. 3GPP Data Rate Evolution**

Following picture shows data rate evolution of 3GPP technologies. In LTE, this evolution is accomplished primarily by adding carrier aggregation and spatially multiplexed layers, using 4x4 MIMO and higher order modulation like 256 QAM. Carrier aggregation allows operators to bond together narrower spectrum from several frequency bands into one larger logical channel. This increases data rate to the user. 4x4 MIMO enables transmission of up to 4 layers of spatially multiplexed streams, using the same air interface time and frequency resources. Compared to single stream, 4x4 MIMO could quadruple effective data rate. Higher order modulation attempts to send more bits of information on a single OFDM symbol, thus enhancing end user bit rate.

5G NR introduces leaner air interface and wider carriers in mid (2.5-6 GHz range) and higher (> 6 GHz range) bands. Data rate enhancement support is available both in base station and UE equipment.

To provide an attractive FWA service, operators should aim to start off higher on the data rate ladder and strive to climb even higher with the right network infrastructure and terminal/CPE solution.

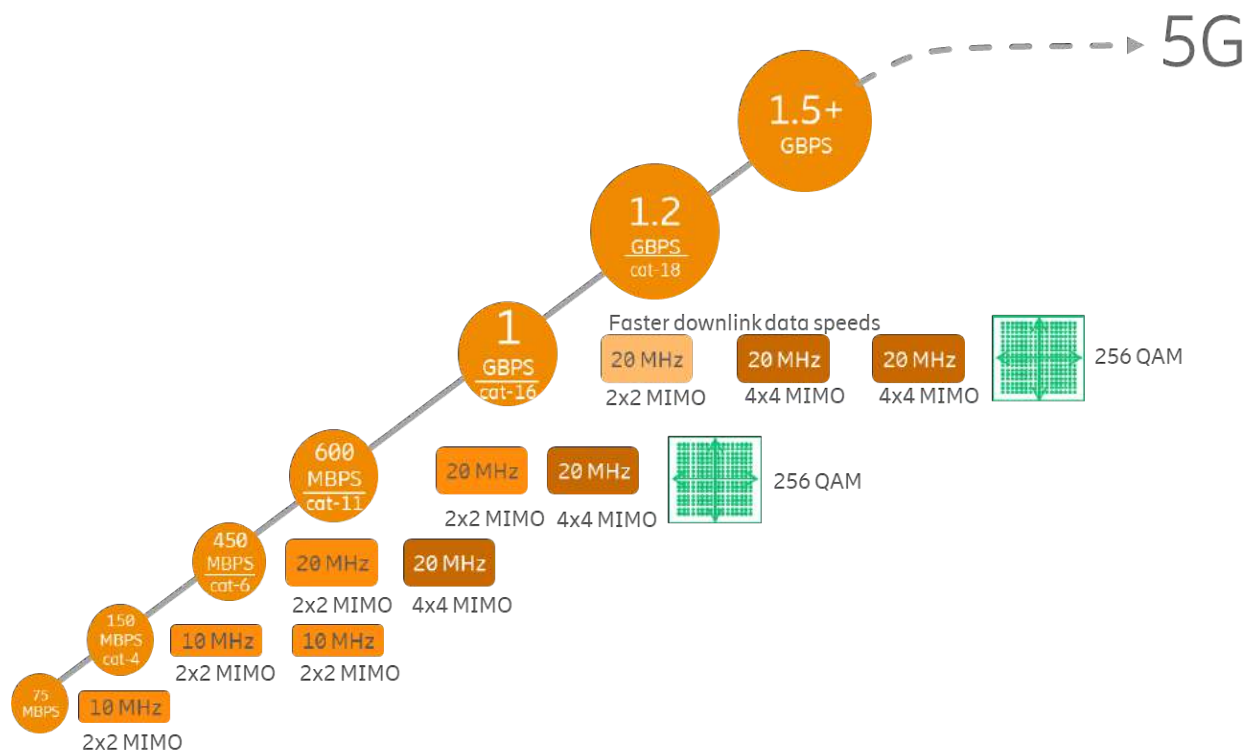



Figure 2: LTE throughput ladder

## 7. Beamforming Technology

Beamforming is the single most important advance in Technology needed for FWA. Beamforming can be done in many different ways. Most of them (not mechanical tilt included) involves time- or phase-shifts of the different antenna elements in an array. Time shift is the most “correct” way to do it, since it is independent of the frequency (e.g. within the bandwidth of the transmitted signal). But phase shift is in most cases easier to achieve practically.


The two main solutions for beamforming are: Digital and Analog Beamforming.

**Digital beamforming** has the benefits that it can be done in the baseband and in the frequency domain. This means that the scheduler and Link Adaptation directly can control the beamforming weights (phase shifts that control antenna pattern) simultaneously for different sub-bands within the same OFDM symbol, as illustrated in the figure below.



**Figure 3: Digital Beamforming**

**Analog beamforming** applies the weights in the time domain, i.e. all symbols within the OFDM symbol get the same weights (same beam direction), which means that Frequency Domain Scheduling, multi-user scheduling or frequency selective beamforming is not possible.



**Figure 4: Analog Beamforming**

## 8. CBRS for FWA

CBRS band in 3.5 GHz spectrum could be a good option for FWA services because:

1. Relatively larger amount of spectrum is available. 150 MHz total allocated for all users. 70 MHz would be dedicated to PAL operation in a licensed mode of operation, while 80 MHz would be available for GAA unlicensed operation. GAA interference and coexistence is managed by SAS. Even though GA is unlicensed, GAA operation is expected to be more predictable because of spectrum coordination functionality provided by SAS.
2. CBRS supports TDD operation, which is well suited to FWA.
3. CBRS provides a good compromise between coverage and capacity. The nature of the spectrum allows practical implementation of massive MIMO radios in this band. Using massive MIMO

radios and multiple user MIMO, capacity enhancement - even over narrower frequency bands – is possible.

4. Strong FWA ecosystem is developing in CBRS band.
5. CBRS implementations are beginning with LTE. We expect migration to 5G in CBRS from 2020.

## 9. Massive MIMO for FWA

Traditional cell sites broadcast signal from LTE sector in all directions covered by that sector. In some places, there would be users that can take advantage of the signal. But most likely, there would be many places where there are no users and the signal broadcasted spatially over wider area would be wasted. Moreover, this wide beam pattern from base station antenna causes inter-cell interference in neighboring cells. The result, from the point of view of a UE, is that total signal levels are lower, and the interference is relatively higher, resulting in a low baseline SINR. Lower SINR implies lower throughput.



**Figure 5: Coverage under broadcast beam**

Massive MIMO aims to solve this problem by employing beamforming. Large number of antennas are used in the base station radio to form narrower beams at the UE location. Because RF power is focused into narrow beams, desired signal level to the UE increases. Additionally, since the cell site is not broadcasting in the coverage area of the sector for traffic channel – it only transmits narrow beams, intercell interference to adjacent cells is reduced. Consequently, SINR experienced by UE improves and so does the throughput. Beamforming tracks channel conditions and strives to maintain optimum beam structure in the cell.


For FWA service, massive MIMO beamforming is useful in enhancing SINR and improving throughputs to the CPE. However, beamforming by itself only provides marginal improvements. Even though we have a higher sector capacity  $X$  in the cell, it is still shared between CPEs served by the cell. If there are 16 CPEs served by the cell at one instance in time, each CPE would get  $X/16$  th of the throughput. These

CPEs are being scheduled in separate resource blocks as indicated by different colors of beams in the picture below, thus sector capacity is split among the CPEs.

Massive MIMO improves capacity but also coverage.

1. SINR improvement over cell edge implies cell edge could be pushed farther from the site, increasing cell range.
2. There is an indirect effect. On many occasions, DL rates are impacted by lack of UL coverage, since UL channel suffering from poor radio conditions may not be able to handle TCP flow control. TCP ACKs/NACKs may not be received from UL channel for DL data transmissions. Large number of base station RX antennas improves UL link budget, extending UL coverage, improving TCP flow control and indirectly helping DL.

5G massive MIMO beamforming adds beamforming to UE in addition to LTE eNB. This could result in further improvement in throughput.



**Figure 6: Beamforming with SU-MIMO**

If the CPEs are orthogonal, so that beams transmitted to the CPEs do not interfere (i.e., overlap spatially), we could allocate the same resource blocks all over again to each CPE. This is shown by the same color of beam in the picture below and is called multiple user MIMO. If the sector capacity is  $X$  and there are 16 (full-buffer) users simultaneously receiving data in the cell, each user can be assigned as much as full capacity  $X$ . There are two benefits of multiple user MIMO.

1. System capacity is enhanced. In the example above, baseline sector capacity was  $X$  with just beamforming. After adding multiple user MIMO, sector capacity becomes  $16X$  (since each of the sixteen users is receiving data with throughput  $X$ ). This is order of magnitude improvement in system capacity, compared to when a single stream was sent to the UE.
2. User throughput improves also. Depending of the SINR experienced by user, their throughput would vary. But since we are allocating potentially all RBs to each user, individual user throughput would improve. For the example above, we assume all users are in good radio conditions, each user's throughput could be as high as  $X$ , which is a 16-time improvement.



We can think of multiple user MIMO creating 16 virtual sectors inside one physical sector in this example.


One consideration for multiple user MIMO is that users need to be orthogonal or spatially separated far enough, such that beams don't overlap. A more intelligent scheduling algorithm would aim to increase UE orthogonality by exploiting not only spatial but also temporal selection of transmissions.



**Figure 7: Beamforming with MU-MIMO**

Since home broadband offers very high data buckets, capacity enhancement of LTE/5G network is always an important consideration. This is true for high subscriber density urban and suburban areas, but also for low density rural areas. In rural areas, improvements in coverage from massive MIMO beamforming may also be of interest. See picture below from field measurements conducted by Ericsson. Here we are comparing DL throughput for 4T4R system with 64T64R systems with massive MIMO. We can see that massive MIMO systems have better range.

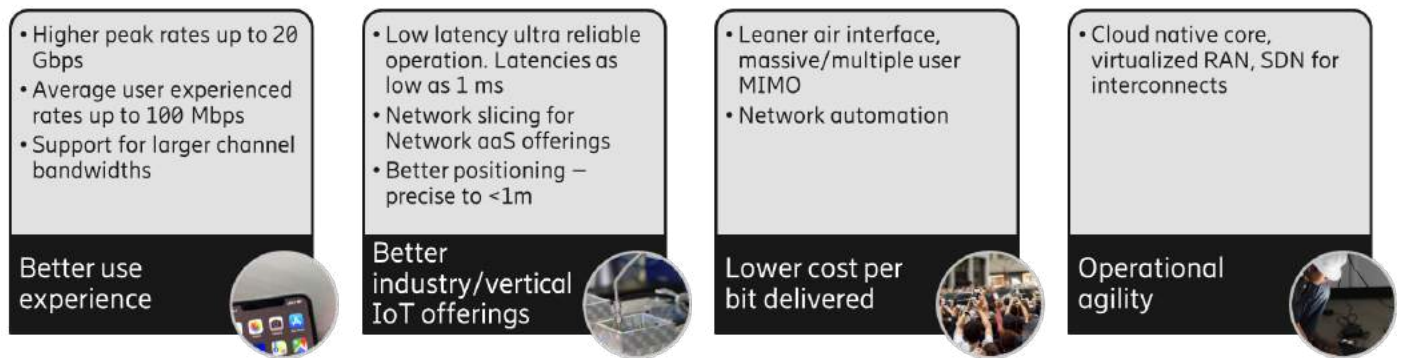




**Figure 8: Comparison of 4T4R and 64T64R user throughput and coverage**

## 10. FWA with 5G

Picture below summarizes key benefits of 5G NR.



FWA operator can utilize each one of these aspects of 5G to their benefit.


1. 5G offers much higher peak and average data rates, allowing FWA operators to offer faster speeds.
2. Converged operators, offering MBB and FWA services, can use 5G's enhanced QoS and network slicing mechanisms to better use common RAN, transport and core infrastructure and provide guaranteed FWA SLAs to their subscribers – reducing their network costs
3. 5G leaner air interface and inherent support for massive MIMO/multiple user MIMO allows operators to transmit more data using limited spectrum, reducing cost per bit delivered. Additionally, mmWave bands with higher channel bandwidths could also allow lower cost per bit of FWA data
4. Cloud native core RAN and transport virtualization, software defined networking and use of COTS hardware enables agile operations and cost-effective networks

## 11. Improving 5G mmWave Reach

mmWave spectrum offers large channel bandwidths and large amount of spectrum. This enables peak data rates up to 20 Gbps. However, propagation of mmWave spectrum is challenging and range is smaller – less than 1 km. FWA deployment of mmWave and mid-band can improve on some of the propagation limitations as below:

1. 5G mmWave antenna arrays can be miniaturized, so large number of antennas can be built in mmWave radios. The antenna arrays allow use of massive MIMO and beamforming. In 5G mmWave, beamforming is used to mostly improve coverage by focusing the RF energy in the direction of intended receiver.

2. Outdoor CPEs become even more critical for 5G mmWave to avoid building penetration losses and to provide line of sight from 5G base station radio to 5G UE.
3. Dual connect operation further enhances coverage, which is typically UL limited. In NR Non-Stand-Alone(NSA) operation, UL from anchor LTE carrier in lower frequency bands could be used to improve cell range. In NR SA operation, a mid band NR carrier can be aggregated with mmWave NR carrier, further improving the effective range of NR cell. This is shown in picture below.




**Figure 9: Combining mmWave and mid band**

#### **Technology of combining mmWave and Mid-bands:**

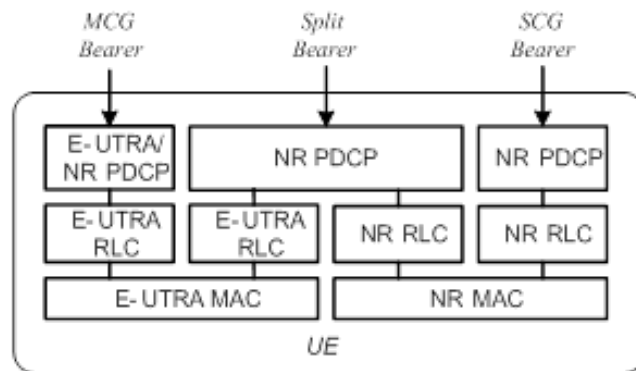
To meet the ever-growing demand for wireless data and applications, service providers must find new sources of spectrum while using existing spectrum efficiently. In addition to enabling new applications and delivering higher data speeds, 5G is unlocking new sources of spectrum in both the mid-band and high-band radio frequencies. Though these new bands are a welcome addition to a service provider's spectrum assets, the mid and high-band coverage is limited by the radio uplink signal quality and characteristics associated with operating within the higher bands. In order to maximize the benefits provided by these new bands, solutions are needed to extend cell coverage.

UL coverage is challenging at higher bands, and also has significant DL throughput impact when losing UL coverage. How much can we extend the range where higher band DL can be used by considering interworking with lower/mid bands?



**Figure 10: Dual Connectivity and Carrier Aggregation**

The first step in the process is Dual Connectivity. In NR NSA operation, UL from anchor LTE carrier in lower frequency bands could be used to improve cell range. Dual Connectivity as defined in 3GPP will allow DL and UL data to switch between LTE and NR independently based on the best performance.



**Figure 11: Dual Connectivity Data Bearer Architecture**

UL and DL decoupling provides enhanced 5G coverage and performance by picking the best NR or LTE leg for uplink and downlink separately. NR coverage is extended by combining the high speed and low latency of NR with the high coverage and high reliability of LTE. Service utilization is maximized because leg switching can be used even with fast changing radio conditions. Allows e.g. the LTE spectrum, with superior coverage, to be used for uplink user data, while the NR spectrum, with superior peak data rate and latency, is used for downlink.

#### Benefits

1. Enhanced NR coverage by leveraging LTE for uplink
2. Increased UL and DL performance

For Dual Connectivity(DC), data is split at PDCP layer and unlike carrier aggregation, Layer1, MAC and RLC will be different for master cell group (MCG) and secondary cell group (SCG). From a link budget perspective, DC will need all L1 control channels on all bands, whereas Carrier Aggregation only needs a limited set of control channels on secondary cells.

With Dual Connectivity, on the UL Data channel (PUSCH) can be moved to low/mid band LTE. The coverage is improved but is often then limited by other UL control signaling.

With Carrier Aggregation, UL L1 control channels can also be moved to NR low/mid band and hence results in a better link budget when combining mmWave NR and mid-band/low band NR.


5G mid-band and high-band coverage limitations by developing a flexible 5G Carrier Aggregation solution which supports control and data traffic on the uplink using a lower frequency band which increases coverage, and on the downlink with a mid or high frequency band which increases capacity and data throughput. With 5G Carrier Aggregation, a service provider can support the 5G uplink operating on a lower band with the 5G downlink operating on a mid, or high-band thereby providing the best of all worlds – better coverage, increased capacity, and higher data speeds.

These solutions enable service providers to optimize the use of their spectrum assets when deploying 5G. A better 5G network will provide more subscribers with higher data speeds while enabling a host of new, low latency applications.

In the scenario considered in the analysis below, aggregating a 5G mid-band with a 5G high-band can improve high-band coverage by up to 7dB extending the high-band cell coverage area by up to 2.5 times. The extended mid-band and high-band coverage also enables a greater offload of traffic from the mid bands to the high-bands providing a higher throughput at the cell edge.


Link Budget Analysis for FWA using mmWave and CBRS using Carrier Aggregation:

- 3.5 GHz, 3X20 MHz, 28 GHz: 400 MHz Carrier Aggregation, CPE : CBRS -> 26 dBm EIRP/20MHz, mmWave CPE : 45 dBm EIRP.
  - Propagation Model : 3GPP RMa LOS, UE height : 5m, RBS Height : 90% Coverage Probability, BS EIRP of 62 dBm on 28 GHz, and 50 dBm/20 MHz on CBRS
1. LOS scenario:



**Figure 12: mmWave and CBRS LOS Bandwidth vs Distance**

2. NLOS scenario:



**Figure 13: mmWave and CBRS NLOS Bandwidth vs Distance**

## 12. Summary of Technical Section

We summarize our discussion of FWA technology below.

1. 3GPP LTE and 5G provide compelling option for FWA services both for existing and new operators.
2. LTE and 5G FWA solutions are applicable to rural, as well as suburban and urban markets.
3. Operators can start off with LTE and then climb the throughput ladder to more system capacity and higher end user rates.
4. Especially for new operators, CBRS provides an attractive way to start up FWA service offerings
5. 5G mmWave spectrum could be useful add on to mid band FWA. 5G mmWave based FWA could also be used in urban areas on its own.
6. Massive MIMO and multiple user MIMO are key techniques to increase system capacity and FWA user throughput.

We conclude that we have the technologies and spectrum today to enable FWA services both for new and existing network operators. In later section of this paper, we show actual FWA field experiences.

### 13. FWA Business Considerations

FWA can help operators with their business challenges:

1. Existing operators that have legacy wireline networks like DSL may find it costly to maintain and to upgrade. Operators that want to expand footprint or upgrade existing capabilities could do that with wireless.
2. The yardstick from broadband has moved since subscribers consume more data and expect faster speeds. Wireless technologies can allow operators to enable fiber like capabilities.
3. In some cases, wireless could be cheaper option than deploying fiber.
4. Expanding wireline plants can be time consuming esp. when it comes to the last mile. Wireless can help operators shorten cycle time and reduce customer churn.
5. Wireless networks are multi-service. Once established for FWA, operators can use them to generate new revenue from services like MBB roaming or IoT connectivity.

### 14. FWA economics

Below we look at factors effecting economics of FWA.

Cost of spectrum. Licensed spectrum could be expensive. For example, operators spent upwards of 40 BUSD for AWS3 spectrum licenses. Since spectrum is expensive, it is important to squeeze maximum utility from it. One strategy is to use inexpensive unlicensed or lightly licensed spectrum for FWA such as CBRS and 5 GHz spectrum. Another strategy is to use mmWave bands (28 GHz, 24 GHz, 37 GHz, 39 GHz etc) for capacity. 37 GHz mmWave also mmWave band offer large amount of spectrum for relatively lower costs.

Cost of network equipment. Since spectrum is scarce and expensive, operators should look at acquiring high performance base station equipment with advanced feature and functionality for FWA like massive MIMO, 5G etc. Similarly, base station sites could be expensive to procure, construct and maintain. A high-performance base station solution that maximizes coverage could also lower the overall costs of deploying and maintaining an FWA network. In previous sections, we saw system capacity improvements

from massive MIMO. Assuming a massive MIMO base station could provide 4x capacity improvement, that will translate into building 4x less sites and spending about 4x less on acquiring spectrum.

Network equipment also includes the LTE or 5G core network and any transmission and aggregation equipment in the operator transport network.

Cost of Site Acquisition and Construction could be significant and even higher than cost of network equipment on the site. Existing operators could leverage their MBB sites to lower site related costs. New FWA operators could also aim to reduce site related costs by deploying high performance base station equipment and by introducing automation in their deployment processes.

Site OPEX comprises of rent payments, backhaul costs, electricity bills and general maintenance costs. Site rent and backhaul comprise the major portion of OPEX costs. TCO for self-owned backhaul would be better than leased line OPEX over the long run. One factor to evaluate here is building of own microwave backhaul.

Cost of CPE. Operators would deploy 100x more CPEs than the number of base stations. For this reason, FWA business cases tend to be sensitive to the cost of CPEs. So, there is a need to drive down the cost of CPE. However, the specs of the CPE could impact network performance and system capacity. A lower spec inexpensive CPE may be detrimental to system capacity – forcing operator to spend more on spectrum and network infrastructure. Outdoor CPEs with high power and high gain antenna enhance network performance and system capacity. The biggest issue with Outdoor CPE is that they require professional installation which dwarfs the cost of the CPE equipment itself. Indoor CPEs could be self-installed by the subscriber. On the other hand, self-install prevents operators from providing SLAs. In the indoor model, you will have opex costs associated with losing signal and hence using more radio resources, churn and in some cases an inability to serve. So outdoor model is best suited for providing a FWA broadband service with SLAs compared to indoor CPE deployments. Operators need to weigh all the cost and performance tradeoffs before crystalizing their CPE strategy.

Technology choice could be critical for the overall business case. A technology that requires frequent upgrades could be costly and may incur disruptions to service. Along this line, effortless upgradability of LTE to 5G (via software without hardware rips) could be beneficial. Since migration of users would not happen overnight, operators should also consider coexistence of LTE and 5G users in the best possible way. Here real time spectrum sharing between LTE and 5G users becomes critical. Without it, operator may be forced to procure new spectrum for 5G.

Subscriber density impact the business case for FWA. A higher subscriber density is preferable to generate higher revenues per cell site and offset operator's network CAPEX and OPEX expenses, but comes at the penalty that higher density makes fixed line competition more cost effective.

Similarly, a higher market share will allow operator to generate more revenue per site up until they reach their spectral capacity. An FWA business case resting on too small or too large market share or density of subscribers could be challenging.

Traffic growth projections factor into the business case in terms cost of future network expansion. This is tied to data speed offerings, since faster speeds may require FWA network densification or more spectrum etc. More traffic and higher data speeds could require more sites, more spectrum or more infrastructure equipment.



## 15. CBRS Business Case Example

It is possible to realize a viable business case for FWA – i.e., an FWA solution with positive cash flow and cost structure better than alternate solutions. Below we show a sample business case for rural and suburban FWA. In this scenario, operator could achieve less than \$300 cost per home passed with a positive cash flow in 3.5 years.

| Budgetary Value*                                                                                                                                                                                                                    | Performance                                                                                                                                                                                                                                                        | Time to Revenue                                                                    |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"><li>– Suburban and Rural:</li><li>– \$&lt;300 Cost per Home Passed</li><li>– 600 homes/sq. mile – Suburban</li><li>– 10 homes/sq. mile – Rural</li><li>– Cash flow positive &lt;3 years</li></ul> | <ul style="list-style-type: none"><li>– 10, 25, 50, and 100 Mbps Service</li><li>– Network supports, 2 Mbps Busy Hour Throughput per HHC with a 28% YoY growth</li><li>– Aligned with Marketing take rates, pricing, and targeted network design metrics</li></ul> | <ul style="list-style-type: none"><li>– Network Deployment can begin now</li></ul> |

## 16. mmWave and CBRS Business Case Example

Below we show a sample business case for rural and suburban FWA using mmWave Frequency and CBRS spectrum. In this scenario, operator could achieve less than \$550 cost per home passed with a positive cash flow in 4 years.

| Budgetary Value*                                                                                                                                                                   | Performance                                                                                                                                                                                                                                                            | Time to Revenue                                                                        |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"><li>– Suburban and Rural</li><li>– \$&lt;550 Cost per Home Passed</li><li>– 500 Homes/sq.mile</li><li>– Cash flow positive &lt;4 years</li></ul> | <ul style="list-style-type: none"><li>– 100, 250, 500, and 1000 Mbps Service</li><li>– Network supports, 6 Mbps Busy Hour Throughput per HHC with a 28% YoY growth</li><li>– Aligned with Marketing take rates, pricing, and targeted network design metrics</li></ul> | <ul style="list-style-type: none"><li>– Network Deployment can begin 1H 2021</li></ul> |

Business case metrics used:

Cost per home passed: This is the cost of building FWA coverage for the number of homes in the service area. This cost does not include customer premises costs.

Cost per home connected: This is the cost of providing connectivity to a customer, including networks and customer premises equipment costs.

Time to revenue: Month or year to cash flow positive Return on investment

## 17. FWA business models


Operators have been trying different business models for FWA. Below, we list a few:

1. Rural high speed FWA. NBN in Australia provides broadband service to rural subscribers. Offering includes 1000 GB bucket with data speed as high as 100 Mbps. To achieve coverage and performance, NBN uses 60 to 100 MHz of licensed TDD spectrum. NBN provides service to otherwise un/underserved communities.
2. 5G FWA. Verizon has launched an FWA service over mmWave 5G variant for \$50-70 per month with no contract. The 5G type FWA service is being deployed in targeted urban and suburban areas and it is meant to compete with existing cable, fiber and other home broadband offerings.
3. Converged MBB and FWA service. Here operator can offer a service to connect a WiFi router, MiFi devices, tablets etc. via the same network. Data consumption from all devices is pooled into a single bucket. Some of the supported devices are mobile and can roam outside of home. This offering has features of both mobile broadband and fixed broadband. T-mobile has just launched the service based on this model.
4. Rural FWA and IoT network: A rural operator provides FWA service and using the same network can offer up low power wide area or massive IoT for smart farms.
5. Rural FWA and MBB roaming: A rural operator deploys FWA LTE network. With network slicing, any residual capacity in the network could be used to accommodate roaming MBB subscribers.

## 18. CBRS FWA Field Experiences

In this section, we look at field trial results from FWA. We highlight performance potential of new technologies like massive MIMO and NR.

In the first instance, we show results from a rural FWA trial using CBRS spectrum. The cell site is configured as 4T4R and we are using two 20 MHz LTE carriers. Outdoor CPEs are used in this set up. Since, this configuration is closer to standard LTE, we can use this as baseline for comparisons.



**Figure 14: Downlink throughput under different RF conditions**

Peak throughput of 180 Mbps in the DL was seen at 0.7 miles from the site. An average home would experience 110 Mbps. One measurement was taken at around 7 miles from the site with 110 Mbps speed in the DL.

Each 20 MHz channel offered average sector capacity of 50 Mbps. Assuming busy hour demand of 2 Mbps per home, each sector would be able to support 50 homes. The number of homes supported by each sector could be increased by adding more CBRS spectrum (more than 2x20 MHz).

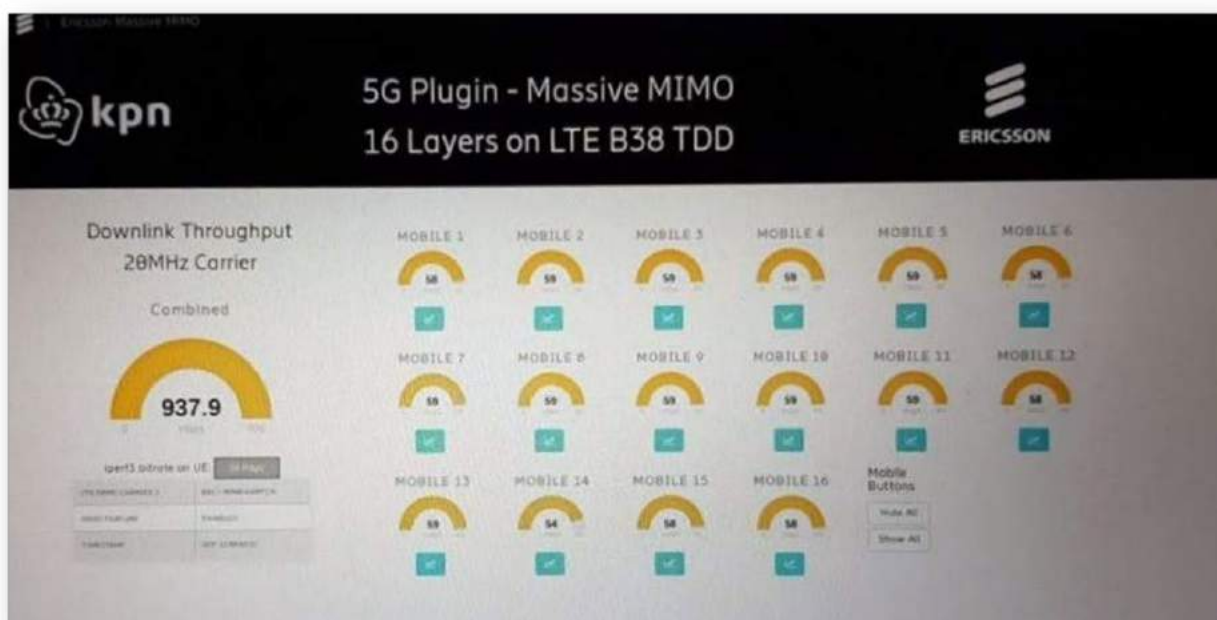
In the next instance, we show how massive MIMO can improve system capacity and individual CPE throughput. We note that CPE placement is ideal, i.e., CPEs are in line of sight and under good RF conditions. They are also spatially separated to minimize inter-beam interference and allow the possibility of achieving good multiple user MIMO gains.



**Figure 15: Setting for multiple user MIMO testing in ideal radio conditions**

These tests use a single 20 MHz LTE carrier with 256 QAM modulation in mid band TDD spectrum. 16 CPEs are configured to transmit simultaneously. System is set up to use up to 16 layers in DL. Total system throughput of ~937 Mbps was observed in the DL.

**Table 3: Results of throughput testing with 16 layer MU-MIMO in good radio conditions**

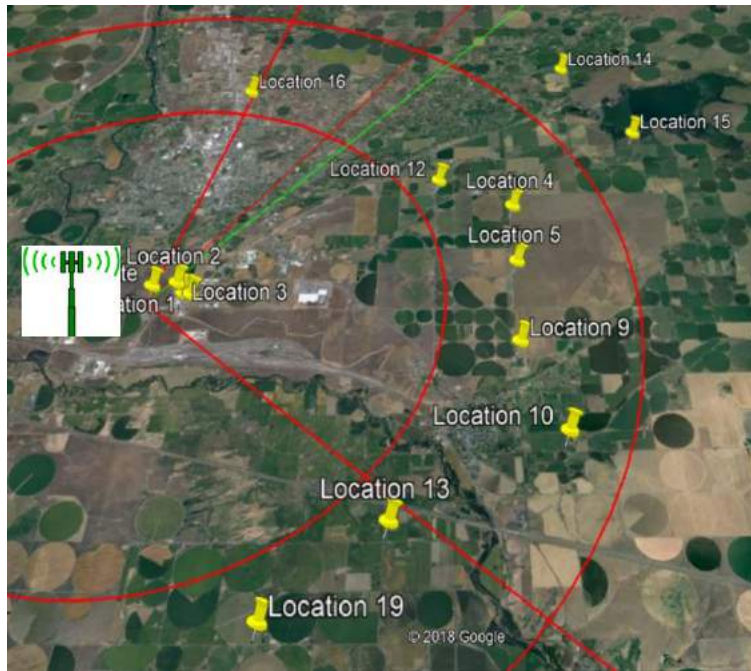


Finally, we show coverage capacity a massive MIMO and 4T4R FWA system in rural area. We used a single 20 MHz carrier in mid band with outdoor CPEs.

With Massive MIMO, high single user throughput performance is achieved with 10 CPEs between 6 to 8 km away from the cell site and 3 CPEs with in 500 meters of the site.

Total system capacity of 250 Mbps was achieved for 8 MU-MIMO layers with 8 CPEs in the 6-8 km distance. The 4T4R, we achieved a system capacity of 80 Mbps in a 20 MHz channel when testing 8 users. So in summary Massive MIMO 64T64R system achieved 3X more capacity compared to 4T4R.

Total system capacity of 400 Mbps was achieved for 13 MU-MIMO layers when tested Massive MIMO product.



**Figure 16: FWA with beamforming and MU-MIMO in rural area**

Also, maximum cell range tests with Massive MIMO products demonstrated that speeds up to 50 Mbps can be achieved in a 20 MHz up to 20 Km distances in LoS conditions.

## **19. mmWave and CBRS Field Trial experiences**

Prototype FWA CPE Devices with mmWave and 3.5 GHz support (Dual Connectivity) are currently being tested. Also, Outdoor FWA CPE devices with mmWave and CBRS will be available 1H 2021.

# **Conclusions**

In this article, we have shown:

1. FWA technology components are in place.
2. There is a viable business case for FWA.
3. Operators have been experimenting with different business models.
4. Deployments that began with rural FWA are moving into suburban and urban areas. Several operators in the US and globally have already been operating profitable FWA business.

FWA has the potential to disrupt existing broadband business. SCTE members could leverage FWA to further build out their own home broadband offerings.

# Abbreviations

|      |                                               |
|------|-----------------------------------------------|
| AP   | access point                                  |
| bps  | bits per second                               |
| FEC  | forward error correction                      |
| HFC  | hybrid fiber-coax                             |
| HD   | high definition                               |
| Hz   | Hertz                                         |
| ISBE | International Society of Broadband Experts    |
| SCTE | Society of Cable Telecommunications Engineers |

# Bibliography & References

<https://www.ericsson.com/en/networks/offerings/fixed-wireless-access>

<https://www.ericsson.com/en/networks/offerings/5g/carrier-aggregation>

<https://www.ericsson.com/en/blog/2020/7/uplink-booster-significant-to-the-new-normal>

Ericsson FWA Handbook. [https://foryou.ericsson.com/FWA-Handbook-June-2019-registration.html?\\_ga=2.209228498.1317989724.1563502297-1427864126.1555962786](https://foryou.ericsson.com/FWA-Handbook-June-2019-registration.html?_ga=2.209228498.1317989724.1563502297-1427864126.1555962786)

# **Roaring Into The '20s With 10G**

A Technical Paper prepared for SCTE•ISBE by

**Dr. Robert Howald**  
Fellow  
Comcast Cable  
1800 Arch St., Phila, PA 19103  
215 286 8037  
robert\_howald@comcast.com

**Robert Thompson**, Comcast

**Sebnem Ozer**, Comcast

**Daniel Rice**, Comcast

**Larry Wolcott**, Comcast

**Dr. Tom Cloonan, Dr. Ruth Cloonan, John Ulm**, CommScope

**Jan Ariesen**, Technetix

# Table of Contents

| Title                                                                | Page Number |
|----------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                 | 5           |
| 2. Scratching the Surface.....                                       | 5           |
| 2.1. Automated MMP: The Profile Management Application (PMA).....    | 7           |
| 2.2. Additional OFDM/OFDMA Specific Benefits.....                    | 9           |
| 3. Executing on the Fiber-Forward Future .....                       | 11          |
| 3.1. Remote PHY Based Distributed Access Architecture (DAA).....     | 12          |
| 4. 10G: DOCSIS 4.0 Full Duplex Progress Update .....                 | 14          |
| 4.1. Key Innovations of FDX.....                                     | 14          |
| 4.1.1. Learn to Share .....                                          | 15          |
| 4.1.2. When You Can't Share, Be Fair .....                           | 16          |
| 4.2. Symmetric Multi-Gigamania .....                                 | 17          |
| 4.3. Network Components to Support FDX .....                         | 19          |
| 4.3.1. Consumer Premises Equipment (CPE).....                        | 19          |
| 4.3.2. DAA Fiber Node.....                                           | 19          |
| 4.3.3. CMTS Core.....                                                | 20          |
| 4.3.4. Home Architecture .....                                       | 21          |
| 4.3.5. Development Progress.....                                     | 21          |
| 4.4. FDX Spectrum Definition .....                                   | 22          |
| 4.4.1. Operator Spectrum Management .....                            | 23          |
| 4.5. More DOCSIS 4.0.....                                            | 25          |
| 4.6. FDX Amplifiers .....                                            | 26          |
| 4.6.1. Echo Cancellation-Based Amplifiers.....                       | 26          |
| 4.6.2. RF Isolation-Based Amplifiers (No EC).....                    | 28          |
| 4.7. Traffic Engineering for FDX Services .....                      | 30          |
| 4.8. Operationalizing FDX.....                                       | 34          |
| 4.8.1. Echo Cancellation and Diagnostics - Live! .....               | 34          |
| 4.8.2. Proactive Network Maintenance .....                           | 37          |
| 4.8.3. DOCSIS 3.0 and DOCSIS 3.1 Participation in FDX Sounding ..... | 40          |
| 4.8.4. "Inverted" Plant of FDX Band .....                            | 40          |
| 4.8.5. Field Tools .....                                             | 41          |
| 5. Bringing Low Latency DOCSIS to Life .....                         | 41          |
| 6. 10G Security Initiatives.....                                     | 44          |
| 7. Conclusion.....                                                   | 45          |
| Abbreviations.....                                                   | 46          |
| Bibliography & References .....                                      | 47          |

## List of Figures

| Title                                                                                                          | Page Number |
|----------------------------------------------------------------------------------------------------------------|-------------|
| Figure 1 - Frequency Split Options Focused on Increasing Upstream Bandwidth .....                              | 6           |
| Figure 2 - Available SNRs Enable More Efficient Modulation Formats and More Capacity with DOCSIS 3.1 [8] ..... | 7           |
| Figure 3 - Time Slicing of the QAM Profiles Associated with Grouping CMs of Like Metrics .....                 | 7           |
| Figure 4 - Modulation Profile Management Platform.....                                                         | 8           |
| Figure 5 - Downstream OFDM Profile Management Operational Dashboard .....                                      | 8           |
| Figure 6 - Upstream D3.0 Modulation Profile Management .....                                                   | 9           |



|                                                                                                                                                        |    |
|--------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| Figure 7 - The OFDM Cyclic Prefix (CP) Guards Against Channel Distortions and Can be Optimized to Channel Characteristics for Maximum Efficiency ..... | 10 |
| Figure 8 - OFDMs Narrow Subcarriers Make It Effective Against Wideband Interference with the Aid of Frequency Interleaving.....                        | 11 |
| Figure 9 - Node Splits vs. N+0 Architecture.....                                                                                                       | 12 |
| Figure 10 - Reduction of Trouble Calls in N+0 (green) v. Traditional HFC (red) Infrastructure.....                                                     | 12 |
| Figure 11 - DAA Aligns Cable Systems to Networking Technology that Leverages Wide Scale Data Center Solutions and Protocols .....                      | 13 |
| Figure 12 - Spectrum Fundamentals of DOCSIS 4.0 Full Duplex (FDX).....                                                                                 | 14 |
| Figure 13 - Echo Cancellation Removes Undesired Transmit Energy at the Receiver.....                                                                   | 15 |
| Figure 14 - Two Key Innovations for FDX: Echo Cancellation and Interference Group Determination and Scheduling .....                                   | 16 |
| Figure 15 - FDX SNR per Interference Group Analysis using Current Tap Performance.....                                                                 | 17 |
| Figure 16 - OFDMA RBAs Example Fully Utilizing the FDX Band.....                                                                                       | 18 |
| Figure 17 - Example of FDX Band Utilization vs Time Enabled by Dynamic Resource Block Assignments (RBAs).....                                          | 18 |
| Figure 18 - Network Component Upgrades for DOCSIS 4.0 Full Duplex.....                                                                                 | 19 |
| Figure 19 - Simplified FDX-Enabled Node Diagram (Single Leg).....                                                                                      | 20 |
| Figure 20 - FDX In-Home Architecture Options .....                                                                                                     | 21 |
| Figure 21 - DOCSIS 4.0 FDX Band Plan [17].....                                                                                                         | 22 |
| Figure 22 - “Classic” View of Mid-Split Based FDX Implementation .....                                                                                 | 23 |
| Figure 23 - DAA Node in FDX Mode with Mid-Split, High Split and FDX CMs – Mid-Split based Diplex... 24                                                 |    |
| Figure 24 - FDX Node Concept as a High Split Augmentation.....                                                                                         | 25 |
| Figure 25 - DOCSIS 4.0 Full Duplex and Extended Spectrum Increase DS OFDM and US OFDMA by Different Means.....                                         | 25 |
| Figure 26 - DOCSIS 4.0 Full Duplex and Extended Spectrum are Complementary in Nature.....                                                              | 26 |
| Figure 27 - An Echo Cancellation-Based Amplifier to Support FDX Beyond N+0.....                                                                        | 27 |
| Figure 28 - Worst Case SNR vs. Echo Cancellation [16] .....                                                                                            | 28 |
| Figure 29 - Evaluating the Signal Relationship of a Diplexer-Less Amplifier (courtesy Technetix) .....                                                 | 29 |
| Figure 30 - Forward and Reverse Freq Response Characteristic of a Diplexer-less Amplifier (Courtesy Technetix).....                                    | 30 |
| Figure 31 - Potential Interference Group Expansion due to Amplifier on an FDX Network.....                                                             | 31 |
| Figure 32 - Breakdown of HSD Traffic Components [4] .....                                                                                              | 32 |
| Figure 33 - Managing FDX Bandwidth to Guarantee Peak Speed Bursts [4].....                                                                             | 32 |
| Figure 34 - Representation of Colliding Simultaneous DS and US Bursts of Shared Spectrum [4] .....                                                     | 33 |
| Figure 35 - Total RF Spectrum Required vs Subs Sharing an IG for 3 Gbps/3 Gbps, 4 Gbps /2 Gbps, and 4 Gbps /4 Gbps [4] .....                           | 34 |
| Figure 36 - Echo Profile from New London, CT FDX Field Trial .....                                                                                     | 35 |
| Figure 37 - CLGD/VSA Echo Profile Coefficient Capture .....                                                                                            | 36 |
| Figure 38 - CATV Plant Map Corresponding to Figure 36 Echo Profile .....                                                                               | 36 |
| Figure 39 - Single Coefficient Variation from Faceplate Removal and Insertion .....                                                                    | 37 |
| Figure 40 - Upstream Adaptive Pre-Equalization Analysis; Echo Response Groups, Mapped.....                                                             | 38 |
| Figure 41 - Downstream Full Band Capture, Standing Wave.....                                                                                           | 39 |
| Figure 42 - Artificial Neural Network, Response Grouping; Self-Organized Map (SOM).....                                                                | 39 |

|                                                                                                                                                                          |    |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| Figure 43 - CDF of RTT Measurements Under Low Utilization Levels for a Mix of Different CM Types and Speed Tiers.....                                                    | 42 |
| Figure 44 - Maximum, Minimum and Mean DS and US LUL with (Left) High Target Latency Values for DS AQM and (Right) Optimized Target Latency Values for DS AQM .....       | 43 |
| Figure 45 - CDF of Connection Latency Based on TCP Handshake RTT Measurements Under Customer Traffic Load for a Mix of Different CM DOCSIS Versions and Speed Tiers..... | 44 |

## List of Tables

| <b>Title</b>                                    | <b>Page Number</b> |
|-------------------------------------------------|--------------------|
| Table 1 - Delay Sources in DOCSIS Networks..... | 42                 |

# 1. Introduction

It seems longer, but 10G unveiled less than 2 years ago! As exciting as the vision is, implementation began well before. 10G has a multi-element roadmap, with its foundations part of MSO initiatives for years. Some are in networks now, while others are in phases of development, proof-of-concept, and trial. Together, they paint a holistic picture of innovation leading to transformative digital experiences for customers.

In this paper, we explore key components of 10G, their synergies, and how they enable new experiences for residential and business consumers. We examine the payoff in speed, capacity, latency, advanced services, reliability and security – all while reducing carbon footprint. In particular, we touch on:

- The symmetric Gigabit foundation of DOCSIS® 3.1 technology
- Advanced DOCSIS 3.1 features – narrow subchannels, shorter cyclic prefixes (CPs), frequency interleaving, diagnostics, and Profile Management Application (PMA)
- Proactive Network Management (PNM) advances in data capture, analysis, localization, and tools for DOCSIS 4.0 technology
- Low Latency DOCSIS capability and implications to gaming and emerging requirements of IoT, M2M, and AR/VR
- Metrics and insights of fiber rich densification at scale, enabling a flexible last mile, increased reliability, lower Opex, and elimination of disruptive node splits
- Experience operating the largest Distributed Access Architecture (DAA) deployment based on remote-PHY (RPHY) and a High Availability (HA) Ethernet switch fabric
- A virtualized core leveraging real-time compute and rapid innovation cycles in the data center space, yielding significant advantages in facility consolidation and costs
- Fiber-rich symmetric networks to power growth in small-to-medium business and enterprise services
- DOCSIS 4.0 features, capabilities, progress, and steps leading to multi-Gig symmetric and full 10G
- The complementary nature of DOCSIS 4.0 Full Duplex (FDX) and Frequency Division Duplexing (formerly Extended Spectrum DOCSIS) FDD

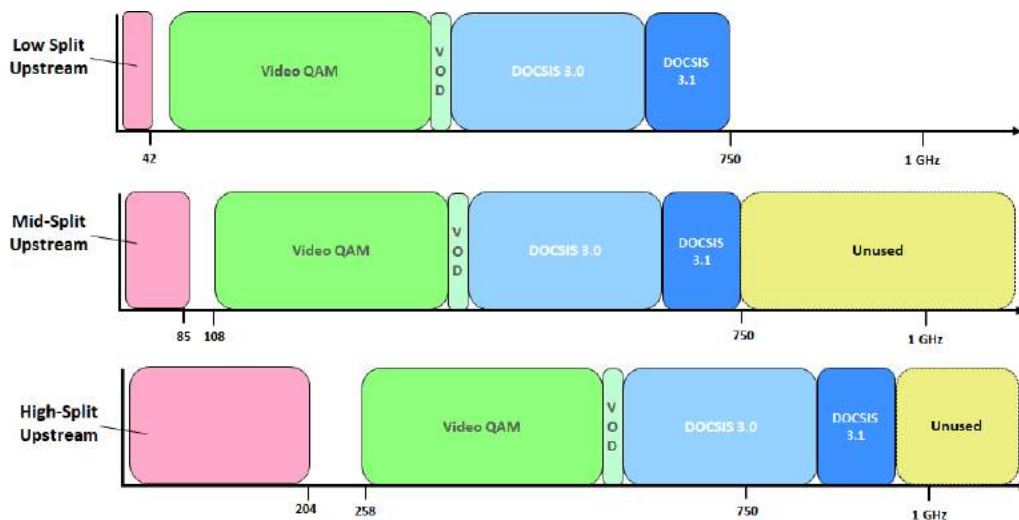
Across Expo Tech Sessions, deeper dives illuminating these and other innovations are explored and described by Comcast subject matter experts focusing on this strategic program.

# 2. Scratching the Surface

It is hard to believe, but the specification development for DOCSIS 3.1 technology began 8 years ago, and service was launched, at Comcast, now over four years ago. Much has been written about the DOCSIS 3.1 foundation changing to OFDM/OFDMA, coupled with a range of higher order modulation formats, including up to 4096-QAM downstream and 1024-QAM upstream, supported by more powerful Low Density Parity Check (LDPC) Forward Error Correction (FEC), and 1 Gbps services.

There are additional features that received less initial fanfare but are receiving more attention now. Figure 1, Figure 2, and Figure 3 highlight these features.


In Figure 1, we compare the common North American upstream band edge of 42 MHz, to the 85 MHz enabled by DOCSIS 3.0 specifications, and to the DOCSIS 3.1 maximum, 204 MHz, often referred to as High Split. The choice of 204 MHz was selected to achieve 1 Gbps of upstream speed and capacity, recognizing that 1024-QAM @10 bps/Hz of raw throughput put that into reach. As such, with DOCSIS 3.1 technology, 1G/1G can be a product offering on the road to 10G.




**Figure 1 - Frequency Split Options Focused on Increasing Upstream Bandwidth**

In Figure 2, the justification for defining Multiple Modulation Profiles (MMPs) is shown. DOCSIS 3.0 is limited to 256-QAM, and encumbered with a less effective FEC. With the combined effect of 1) better FEC (a lower SNR required for a given QAM format), 2) improved plant fidelity over time with deeper fiber and shorter amplifier cascades, and 3) migration to DAA, Cable Modems (CMs) are capable of much better bandwidth efficiency than 256-QAM. The dB attributable to the above can be used for capacity gain rather than just better performance margin of 256-QAM. Figure 2 shows why this is effective – many CMs that can support higher QAM efficiency than 256-QAM.

Lastly, note that it was determined during this study that this national SNR distribution also exists down to the service group (SG) level. Because of this, it makes sense for there to be different modulation profiles defined within a single SG, and these profiles cycled through based on their traffic demands on the time axis. This is shown in Figure 3, where sets of devices (color-coded) that exhibit common MER metrics have a time and a frequency slice of the downstream broadcast carrying their traffic.



**Figure 2 - Available SNRs Enable More Efficient Modulation Formats and More Capacity with DOCSIS 3.1 [8]**

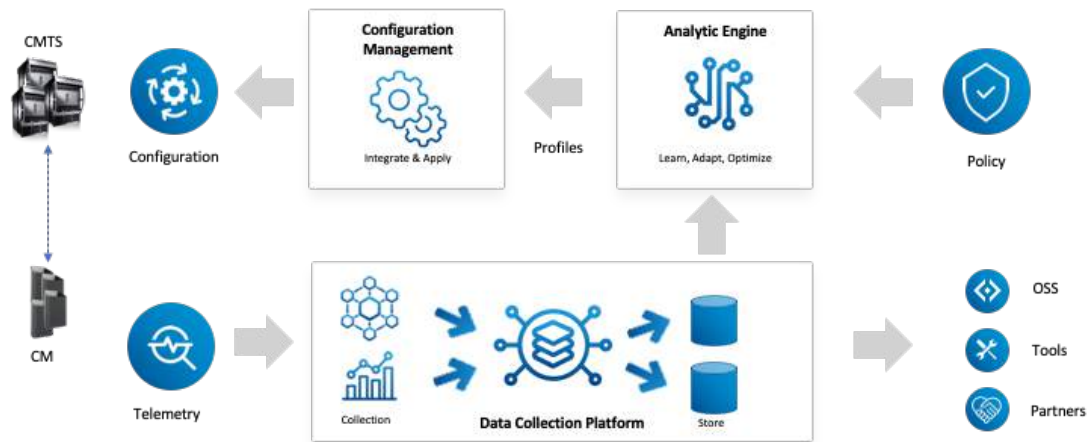


**Figure 3 - Time Slicing of the QAM Profiles Associated with Grouping CMs of Like Metrics**

## 2.1. Automated MMP: The Profile Management Application (PMA)

Efficiently managing the modulation profiles of DOCSIS 3.1 has proven to be a key component in successfully achieving the capacity promises of the technology. As DOCSIS 3.1 was deployed, it quickly became clear that dynamic and autonomous profile management is key to a holistic capacity plan [13]. Similar technology for DOCSIS 3.0 was recently deployed during the COVID pandemic to increase upstream capacity [14].

A PMA system leverages network data, such as RxMER per subcarrier in downstream OFDM channels, along with traffic, codeword, and signal level statistics to make optimization decisions that increase the capacity of the downstream, as shown in Figure 4.



**Figure 4 - Modulation Profile Management Platform**

This same cloud platform has been used to optimize the capacity of DOCSIS 3.0 upstream channels. The optimization is based on three factors: upstream SNR, correctable and uncorrectable codeword errors, and signal levels. The metrics can be used by the analytics engine to adjust FEC and modulation for the data, Unsolicited Grant Services (UGS), and station maintenance Interval Usage Codes (IUCs). For DOCSIS 3.1 upstream transmissions, as in the downstream, optimization can be done by the analytics engine, which designs different modulation levels across the spectrum, tailored to channel fidelity.

This platform has enabled us to increase our access network capacity by over 36% in the downstream signal direction, and 20% in the upstream. The 36% is relative to a 256-QAM baseline and means PMA has brought the average efficiency close to 2048-QAM (37.5%). There is significant untapped gain in the upstream at this point, since 20% is comparing only different DOCSIS 3.0 settings. These results are tracked in operations as shown in Figure 5 and Figure 6.



**Figure 5 - Downstream OFDM Profile Management Operational Dashboard**



**Figure 6 - Upstream D3.0 Modulation Profile Management**

Beyond capacity gain for existing channels, this solution also enables spectrum to be used that was not previously possible. In some cases, spectrum is being used for OFDM downstream transmissions that are 100 MHz beyond common HFC frequency boundaries, driven by the pandemic-driven increases in traffic. In the upstream, a 6th DOCSIS 3.0 channel is being added below 15 MHz -- a notoriously noisy region that is suddenly more plausible when a PMA tool can be put to work to select the best FEC and modulation mix for a quality customer experience.

These PMA tools also help to ensure network robustness. For example, in the downstream spectrum, interference from mobile 4G Long Term Evolution (LTE) carriers is known to cause instability and loss of capacity without a system to effectively manage it. In the upstream, a variety of different noise sources [10] can be mitigated by a modulation profile management system. While these solutions mitigate the impact of the noise while maximizing the customer experience and capacity, the actual noise itself remains. It's still important to dispatch technicians to remediate the underlying issues. Combining profile management with PNM provides the best of both worlds. Technicians can be deployed to a specific network location that needs attention, and the customer is completely un-impacted and unaware through it all.

Similar techniques are being developed for upstream OFDMA channels, to optimize modulation profiles for challenging spectrum areas. The upstream signal path is a skinny, scrappy portion of total available capacity, routinely trampled by interference from FM radio, cascaded filter distortions, or the lower frequency noise common today, not to mention spurious/impulse noise. For DOCSIS 3.1, not only the modulation, but the pilot pattern, CP, and Roll-off Period (RP), along with subcarrier and mini-slot frame size are important attributes to optimize for robustness and capacity. These are important, for example, to achieving 1Gbps speeds in a high split configuration.

As product requirements grow beyond symmetrical Gbps to multi-Gbps symmetrical services, DOCSIS 4.0 Full Duplex (FDX) will be used. This technology will also benefit from OFDMA and OFDM modulation profile management. FDX introduces some new concepts that can also be managed as virtual network functions to optimize use of a shared spectrum resource, including Interference Group (IG) detection and Resource Block Assignment (RBA) usage.

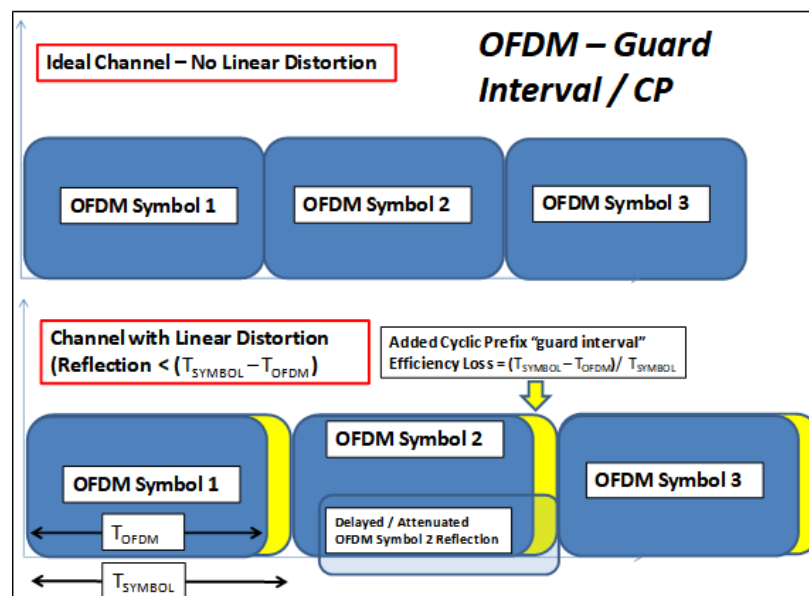
## 2.2. Additional OFDM/OFDMA Specific Benefits

Two other features of DOCSIS 3.1 are worth mentioning. A key mechanism of how OFDM signals deal with linear distortion that results in Intersymbol Interference (ISI) is through a simple delay between OFDM transmission symbols. For frequency response distortion, it is also a simple mechanism – a single



complex multiplier applied to each subcarrier. Compared to the complexity of tapped-delay line SC-QAM adaptive equalizers for linear distortion compensation, these are simplifications that come directly from the OFDM signal structure.

The time delay inserted is determined by understanding the delay spread of the channel, and the idea shown in Figure 7. The delay is known as the “guard interval” and is implemented by the addition of a CP of redundant information that is easily generated as part of the IFFT algorithm that creates the OFDM symbol. The CP is wasteful overhead, but necessary for clean transmission – the SC-QAM parallel being the spectrum shaping factor “alpha” that amounts to bandwidth overhead beyond the minimum Nyquist bandwidth to transmit a QAM symbol of a particular symbol rate.

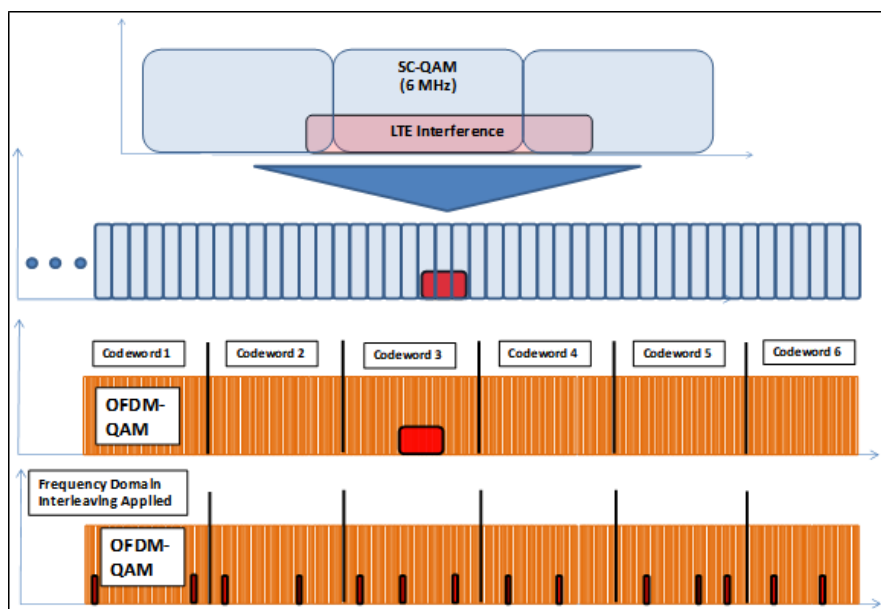


**Figure 7 - The OFDM Cyclic Prefix (CP) Guards Against Channel Distortions and Can be Optimized to Channel Characteristics for Maximum Efficiency**

Another DOCSIS 3.1 feature possible only because of OFDM is the frequency interleaver function shown in Figure 8. Each subcarrier is carrying a QAM symbol; a set of subcarriers form a codeword. However, when frequency domain interference strikes, only some of the OFDM subcarriers may be impacted. It is an OFDM advantage that subcarriers without interference continue to carry capacity, whereby an SC-QAM signal may be wiped out completely.

What the frequency interleaving attribute offers is the ability to spread the interference out such that it is not concentrated into a contiguous set of subcarriers or in a single codeword. By shuffling subcarriers prior to transmission, and then un-shuffling at the receiver, the tones with interference can be spread across the entire allocated OFDM band. They will then fall across multiple codewords, whereby each codeword will now have less interference than the single impacted codeword without interleaving, improving overall error correction performance.





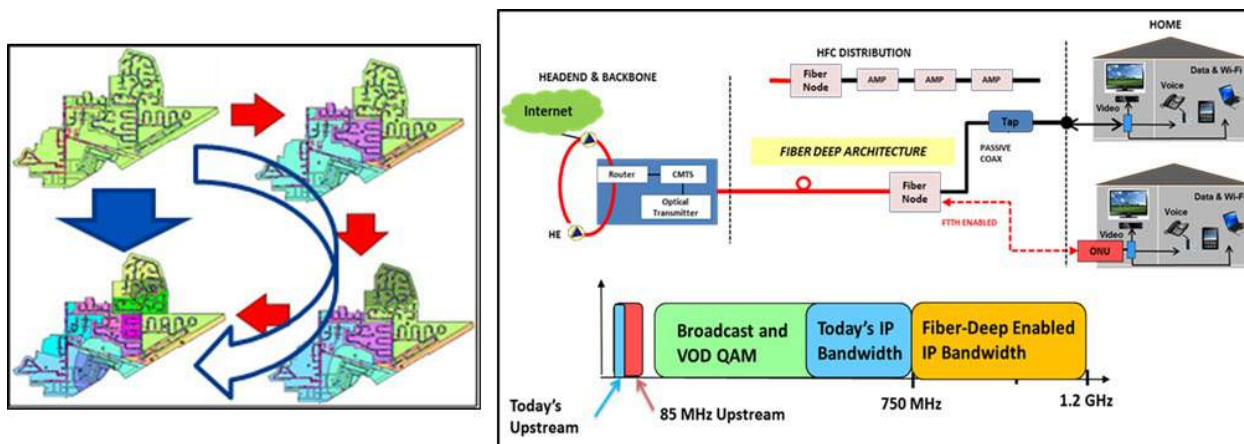
**Figure 8 - OFDMs Narrow Subcarriers Make It Effective Against Wideband Interference with the Aid of Frequency Interleaving**

### 3. Executing on the Fiber-Forward Future

Cable operators have been pulling fiber deeper into their network and splitting nodes for years, as one tool to manage the Compounded Annual Growth Rate (CAGR) of bandwidth usage and consequent capacity expansions. Node splitting reduces the number of customers in a service group by approximately two, doubling average capacity available to each home. Adding or reallocating channels of existing spectrum are other key tools to address CAGR. Spectrum flexibility, however, is very limited in today's upstream.


A “node split” can refer to a physical split, or the segmenting of an existing node by populating it with additional modules supported by new CMTS ports in the Headend. As nodes continue to be split and the serving group size shrinks, the efficiency of these splits decreases and the cost increases, because there are more physical splits rather than segmentation. A more forward-looking approach is to split multiple nodes at once. When many splits are projected over several years in an area, it may make sense to perform pre-emptive splits for labor efficiency and get ahead of trends as opposed to revisiting the same area repeatedly.

A version of “multiple node splits at once” is to be more methodical about the end architecture, as the node splitting occurs in these high growth areas. An approach we've taken in such areas replaces haphazard node splits with disciplined designs that optimize HFC design around a node with zero trailing RF amplifiers (N+0) end state, shown in Figure 9. Millions of homes passed are connected via this method today, and these networks will not have to endure disruptive node splits for many years, and possibly never.



**Figure 9 - Node Splits vs. N+0 Architecture**

The N+0 architecture provides long-term capacity, improves End of Line fidelity, increases reliability, and adds substantial new spectrum currently limited by RF amplifiers in traditional HFC topologies. Operational performance is carefully monitored. Figure 10 is an example of comparative trends in Trouble Call (TC) performance between N+0 and standard HFC networks in a Comcast Division actively building N+0. Furthermore, the economics are favorable when compared to repeated node splits over time, especially in high growth areas of the network [7].



**Figure 10 - Reduction of Trouble Calls in N+0 (green) v. Traditional HFC (red) Infrastructure**

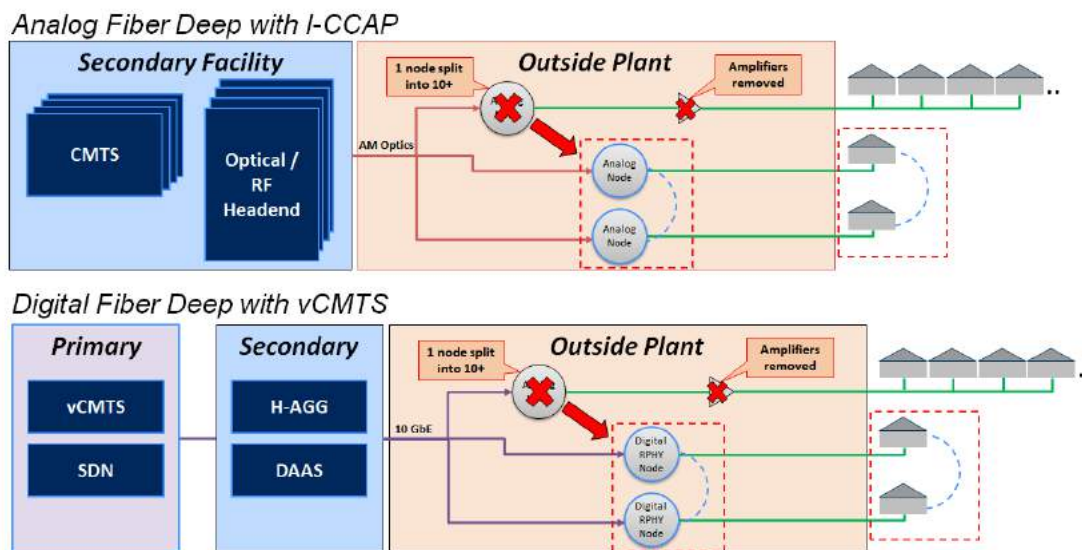
### 3.1. Remote PHY Based Distributed Access Architecture (DAA)

Traditional HFC networks have long used cable-specific analog optics to carry signals downstream; for the upstream, either similar analog optics or proprietary “digital return” systems. Our DAA deployments are based on Remote PHY device (RPD) technology and include all new N+0 systems because of the massive scale efficiencies, among other benefits, to the network and facilities. RPHY is by far the widest DAA flavor deployed. A brief summary of DAA benefits:

- Improved reach and wavelength efficiency of digital fiber versus analog
- Fidelity gains (MER) due to the removal of analog optics
- Physical scalability in the inside plant with the removal of RF cabling and combining networks
- Operational simplification, network availability, cost reduction
- Introduction of the global Ethernet ecosystem into the optical access network
- The vision of a network agnostic (at last!) to last-mile access technology
- Path to CMTS virtualization (vCMTS)

Regarding the last bullet, with DAA, a major part of the DOCSIS-specific CMTS functionality is moved into the plant. The subset that remains – packet processing, switching, storage, scheduling -- can be revisited from the perspective of today's compute power and the capabilities of real-time software. Commercial off-the-shelf (COTS) servers and switches can be combined with software to implement the CMTS function virtually. Our first DAA deployment, hosted by a virtual CMTS, or vCMTS, was introduced in 2018.

Figure 11 compares an N+0 network based on traditional AM optics and I-CCAPs to a DAA-based vCMTS solution. In the diagram, DAAS stands for Distributed Access Architecture Switch, and H-AGG stands for Hub Aggregation (also switching / concentration). Note the shifting of the CMTS core location to be further northbound in the vCMTS architecture, owing to the capabilities of Ethernet optics in terms of reach and wavelengths. This leads to opportunities for facilities reductions and consolidations in a massive way.



**Figure 11 - DAA Aligns Cable Systems to Networking Technology that Leverages Wide Scale Data Center Solutions and Protocols**

## 4. 10G: DOCSIS 4.0 Full Duplex Progress Update

DOCSIS 4.0 Full Duplex (FDX) enables multi-gigabit symmetrical capability over a coaxial last mile, a massive increase of upstream capacity and speeds. With the nature of consumer traffic being highly asymmetrical, cable operators optimized broadband for this reality, but are now getting ahead of requirements for new services and emerging applications that tax the upstream path – high definition (HD) home security cameras, cloud gaming, IoT, machine-to-machine applications, and those we cannot yet imagine. What FDX achieves is significantly more upstream bandwidth without a downstream bandwidth penalty, delivering new speeds while staying on top of CAGR of data traffic. Another key 10G pillar is low latency, which is also a key requirement of future applications, including many of those described above.

### 4.1. Key Innovations of FDX

Figure 12 shows the fundamental concept of FDX – access to massively more upstream. The FDX upstream is based on the same DOCSIS 3.1 physical layer – 96 MHz OFDMA blocks of 25 kHz or 50 kHz subcarriers, the same QAM profiles, and the same FEC. Six possible 96 MHz OFDMA blocks are added in the 108-684 MHz band as a complement to an 85 MHz “mid-split” system.

Figure 12 also makes clear the fundamental challenge of FDX – the new upstream overlaps with the downstream. This differs from the basic and, until now, exclusive Frequency Division Duplex (FDD) principle of cable TV. This dates back to the 1970s, when downstream and upstream signals began to exist in separate regions, as in Figure 1. Coaxial cable is an inherently broadband medium, necessary originally to handle a multiplex of analog video channels, which deferred the need to be especially efficient with its use.

Meanwhile, phone companies had to work with much poorer physical medium – twisted pair copper wires -- to enable data services on a system optimized for voice bandwidth. Later, wireless networks had to manage within regulated spectrum. In both cases, spectrum efficiency was at a premium from the outset. Consequently, powerful technologies, including a technique known as Echo Cancellation (EC), were developed and matured to use available bandwidth in full-duplex mode. Years later, cable systems are now poised to leverage multi-octave, broadband EC as one of two key innovations of FDX, as efficient use of cable spectrum becomes increasingly important.

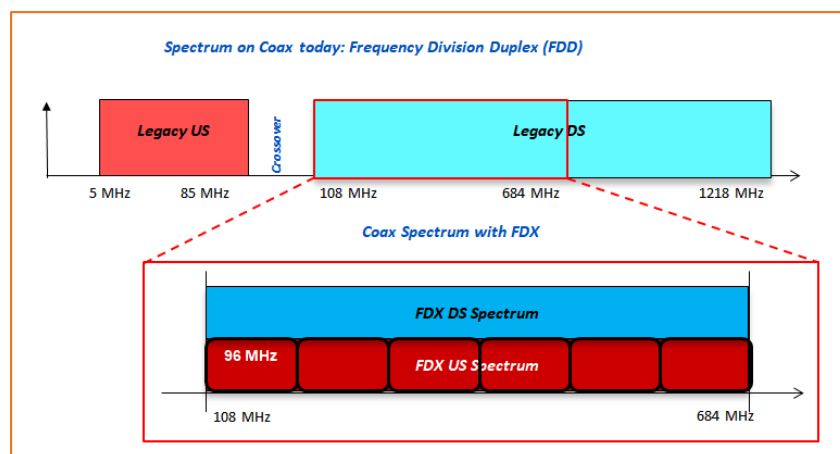



Figure 12 - Spectrum Fundamentals of DOCSIS 4.0 Full Duplex (FDX)

#### 4.1.1. Learn to Share

The unique aspect of FDX is that the 108-684 MHz upstream can also be occupied by DOCSIS 3.1 downstream traffic. Echo Cancellation allow the simultaneous use of spectrum by removing the undesired signal from the desired signal at the receiver when they are sharing common spectrum. The term “Echo Cancellation” captures that: from a DAA node receiver perspective, it operates on undesired reflections of the downstream that bounce back from the plant. In fact, some of strongest interfering energy originates within the node itself, as either limited RF isolation or an internal reflection. Regardless of origin, it is handled by the same EC mechanism.

From the FDX RPD perspective, the EC function is conceptually straightforward – the downstream must be “subtracted” before the upstream receiver can process the upstream data. This is achieved through RF isolation of the downstream signal from the upstream receiver, and the ability to cancel residual copies of the downstream that make it to the receiver. It relies on the fact that the receiver “knows” the signal being transmitted and can therefore theoretically subtract the interference, as shown in Figure 13.

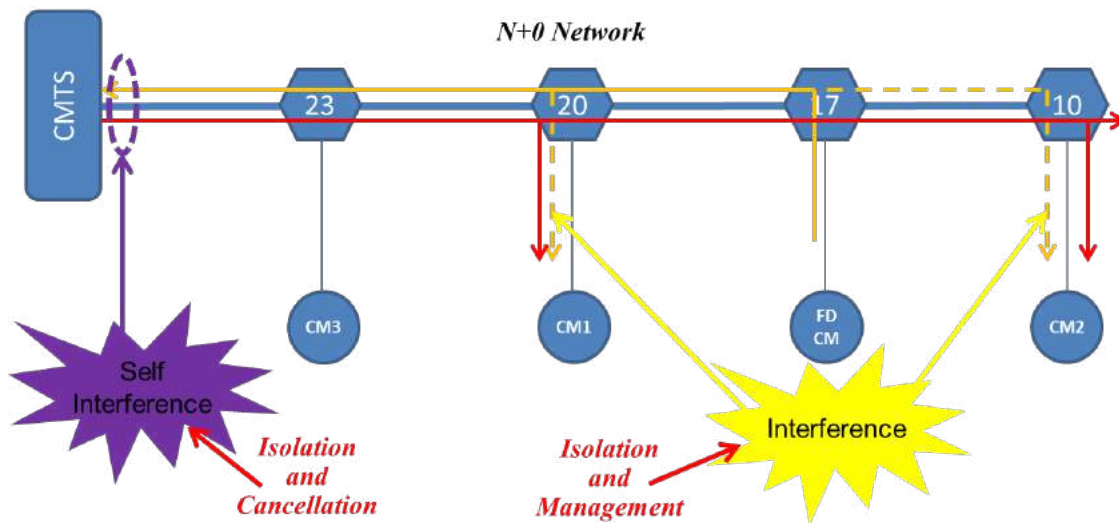


**Figure 13 - Echo Cancellation Removes Undesired Transmit Energy at the Receiver**

In practice, the upstream receiver must also measure the noise from the downstream transmitter and cancel it out for highest fidelity and maximum upstream bandwidth efficiency.

One important architectural difference of HFC compared to the telco last mile is that HFC is point-to-multipoint. This difference means one additional innovation is needed for FDX. That innovation involves clustering the users into groups that could inadvertently interfere with one another if any in the group were listening to the downstream while another in the group was transmitting. This is an extra scheduling exercise for the CMTS, which first determines the clustering, and then as part of its normal job of assigning transmit time slots to modems checks the upstream bandwidth allocation map (MAP) against downstream destinations, and if necessary, adjusts to avoid possible collisions.

These clusters, an important element of FDX technology definition and management, are called Interference Groups (IGs) and aggregates of IGs are called Transmission Groups (TGs). Figure 14 represents the two new techniques required to make FDX work, from the perspective of the FDX RPD.



**Figure 14 - Two Key Innovations for FDX: Echo Cancellation and Interference Group Determination and Scheduling**

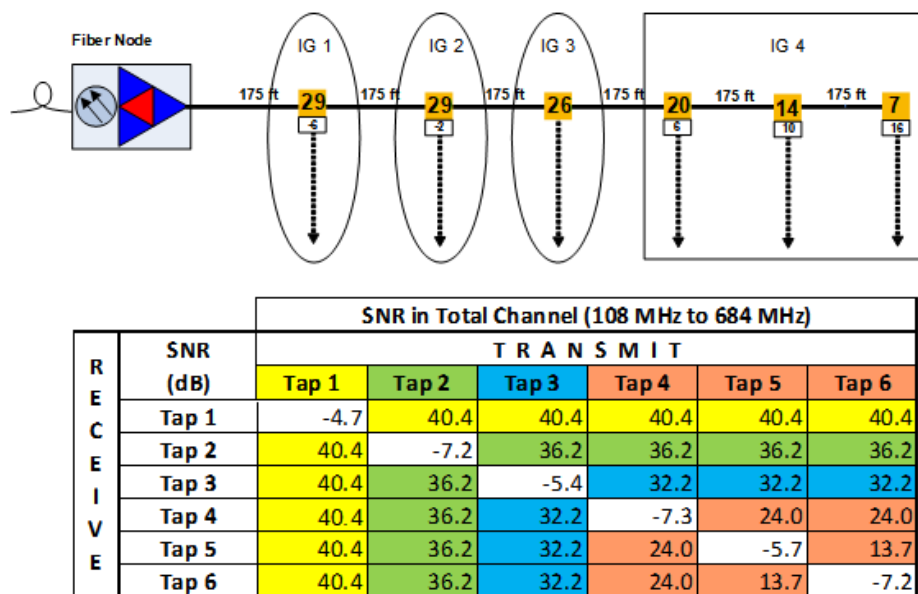
#### **4.1.2. When You Can't Share, Be Fair**

Echo Cancellation is based on having or gaining knowledge of the transmit signal and noise to be cancelled. This can be a challenge in a point-to-multipoint architecture, and from the CPE's perspective. While an FDX CM has knowledge of its own signal, it will not have knowledge of its neighbor's. Only enough neighbor-to-neighbor RF isolation will prevent an FDX band transmission from interfering with another's downstream receive signal.

Unfortunately, enough isolation within a single tap or across taps cannot always be guaranteed. Some taps, in particular mid-to-high value Taps (29 dB, 26 dB, 23 dB), will have sufficient RF isolation from tap to tap. These dB relationships will be calculated by the CMTS as part of the network "sounding" process, used to determine IGs, TGs, and QAM profiles, and update them periodically. Figure 15 is an analysis of a single RF leg in an N+0 environment of a likely IG/TG determination using typical tap specifications.

As shown in Figure 15, users on the same tap are naturally going to form an IG. However, as we will see later, the statistical likelihood of such a small group competing for use of FDX bandwidth in a way that impacts services is negligibly low until FDX penetration becomes very high (for which further features can be enabled). This leads to defining TGs – aggregating IGs simplifies the scheduler aspect by having a smaller number of groupings to manage: Divide to conquer.





**Figure 15 - FDX SNR per Interference Group Analysis using Current Tap Performance**


Figure 15 also shows the FDX band SNRs for each IG. Beyond the “same tap” limitations (shown in the white cells, with negative SNR if an FDX CM is transmitting at the same time), we can see, for example, within IG4, a useable but low SNR would lead to 128-QAM maximum and as low as 8-QAM when an FDX CM on one of these taps was transmitting. By forming an IG across these taps, these low SNRs are avoided by choosing different downstream and upstream time windows for the users in the IG.

Once IGs are established, the CMTS scheduler executes its “normal” operation with the additional rules. The CMTS scheduler is the arbiter of access to the network, allocating time slots for downstream and upstream transmissions to fairly and efficiently serve the aggregate demand. The CMTS does not typically or explicitly need to pay attention to a single user’s downstream and upstream access. That changes with FDX, as there is an awareness needed that accounts for IG/TG relationships. This extra scheduler “step” makes sure that a downstream receiver signal of a potentially vulnerable user is not interfered with by an adjacent FDX upstream.

## 4.2. Symmetric Multi-Gigamania


What happens as FDX users increase in penetration? To enable fair and efficient access, these IGs and TGs can be granted slices into sub-blocks of the very sliceable OFDMA spectrum, so that availability of FDX bandwidth is always possible for any TG.

Figure 16 shows an example whereby the entire 108-684MHz FDX band is activated. Three TGs each have 384 MHz of downstream and 192 MHz of upstream capacity in the FDX band, or about 3 Gbps downstream and 1.5 Gbps upstream (including the 85 MHz legacy upstream). Sub-band allocations are called Resource Block Assignments, or RBAs. Of course, the FDX band is not all of the available DOCSIS downstream. For example, there could be another 192 MHz of non-FDX D3.1 spectrum in the downstream, for a total of 5Gbps, and also DOCSIS 3.0 spectrum.



**Figure 16 - OFDMA RBAs Example Fully Utilizing the FDX Band**

FDX requirements have been written such that RBAs can be dynamically allocated as traffic demands shift with time. Figure 17 demonstrates this feature, with FDX band allocation on the vertical axis and time on the horizontal axis.



**Figure 17 - Example of FDX Band Utilization vs Time Enabled by Dynamic Resource Block Assignments (RBAs)**

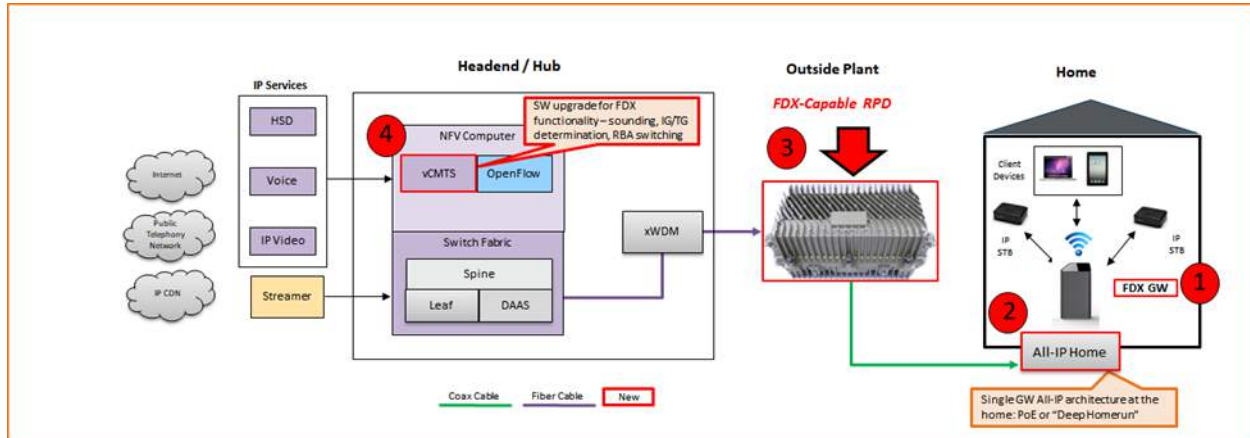
Two last important points are in order with regard to FDX spectrum:

- 1) Because FDX is based on DOCSIS 3.1 technology, FDX can only be activated as FDX spectrum where downstream DOCSIS 3.1 spectrum is located.
- 2) Existing DOCSIS 3.1 devices using the FDX band must be aware of FDX, as this ensures they are effective participants in the IG/TG mechanism. This mode – a software (SW) upgrade to existing CPE – is known as FDX-Light (FDX-L) mode.



### 4.3. Network Components to Support FDX

Figure 18 identifies components of the end-to-end network effected by FDX. The red highlights identify areas of the network impacted by FDX that we will describe in the sections below.



**Figure 18 - Network Component Upgrades for DOCSIS 4.0 Full Duplex**

#### 4.3.1. Consumer Premises Equipment (CPE)

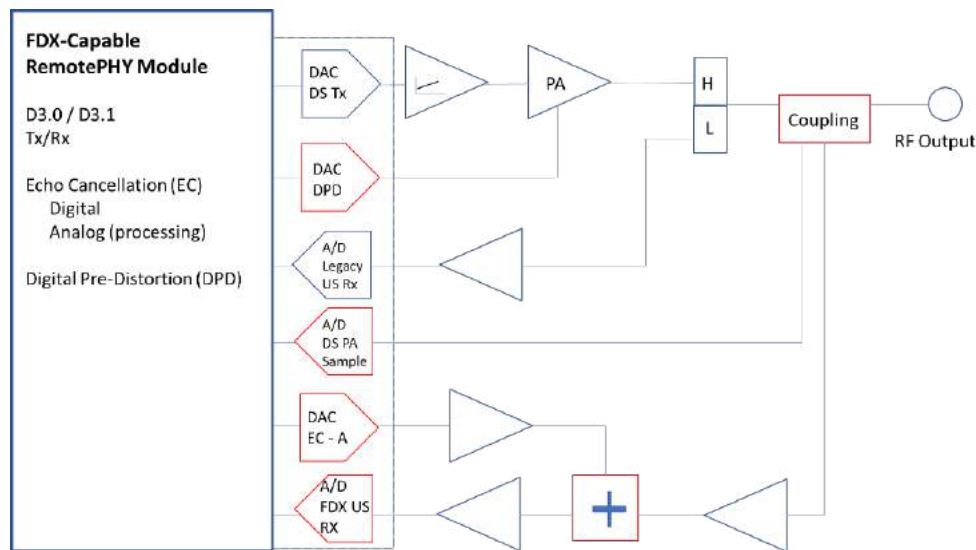
As with DOCSIS 2.0, DOCSIS 3.0, and DOCSIS 3.1, support for FDX requires new silicon development to implement the wider upstream, the additional OFDMA blocks to fill it, and to enable EC. Support for the new scheduling aspects are also necessary, although this is SW functionality and not necessarily a new hardware requirement. Existing DOCSIS 3.1 modems are expected to participate in an FDX system in “FDX-L” mode, as previously described.

The overlapping spectrum means that the CPE gateway RF design deviates from the typical diplexer-first design, relying on an arrangement based on splitters/couplers to accommodate the overlap while maintaining legacy bands.

#### 4.3.2. DAA Fiber Node

Following the CPE upstream to the node, any node in an FDX system must also support the expanded bandwidth, OFDMA blocks, and Echo Cancellation. As with the home gateway, the internal RF design of the node must account for overlapping bandwidth in a way different than a traditional FDD node. This is one of the more complex parts of the node design – how to manage downstream and upstream paths to support additional couplers, while integrating EC functionality and minimizing RF losses affecting levels.

An example block diagram is shown in Figure 19, with new FDX items highlighted in red. The diagram represents a single node port or RF distribution and receive leg.



**Figure 19 - Simplified FDX-Enabled Node Diagram (Single Leg)**

It is possible that a DAA node can be a traditional FDD node on Day 1, and an FDX node later, when needed. Such a node would become FDX via a software upgrade and appropriately built-in hardware processing and RF capability. This would prevent future visits for FDX hardware upgrades.

#### **4.3.3. CMTS Core**

As indicated, CMTS “sounding” determines IG/TG relationships. Fundamentally, the sounding process measures relative relationships of devices to one another to determine who may need to be kept apart to avoid interference when spectrum is being shared. It uses that information to adapt the scheduler to prevent interference when FDX CMs are transmitting. It does this by scheduling them in time slots that keep them away from vulnerable receivers based on the IG/TG matrix.

Note that IG/TG relationships are not static. New users can be added and subtracted, and devices may be moved around in a home or business, affecting network characteristics. While not very dynamic, the system has to periodically re-execute “sounding” to establish the correct IG/TG matrix for the scheduler.


The network itself is also not static. Again, while it does not change dynamically, it can move slowly over time and temperature, and/or have maintenance done that affects channel characteristics. The EC must track these and adapt accordingly. There are not explicit transient behaviors or convergence time requirements in the specifications, so this is an area that has seen significant characterization and is ripe for potential optimization and differentiation.

Lastly, as shown Figure 17, in the feature-complete implementation, the scheduler will support optimized RBA allocations that follow the needs of the traffic demands and service tier distribution among the IGs. Today’s DOCSIS 3.1 schedulers are sophisticated, and their capabilities not fully leveraged. FDX will add the additional layer of IG/TG no-go zones – at first as simple as a “look-up table,” but that will eventually become more dynamic based on traffic patterns, as use of the FDX band becomes better understood.

#### 4.3.4. Home Architecture

The FDX specifications are written from a system engineering standpoint, assuming a point-of-entry terminating home gateway. Analysis was done that considers a single gateway termination, but one that may be located deep in the home, increasing RF loss. We refer to this as the “deep homerun.” The advantage of this approach would be that it would give the customer the ability to center the device for Wi-Fi optimization. Also, importantly, a self-install kit (SIK) approach can be considered as a cost-effective way to deploy new devices. SIKs are also highly preferred by customers, compared to visits from technicians, pandemic or no pandemic.

These home install options are shown in Figure 20.



**Figure 20 - FDX In-Home Architecture Options**

With the practical complexities of having FDX upstream signals within the legacy downstream bandwidth of STBs, it is expected that an FDX home will become an All-IP home. This makes the “deep homerun” SIK approach more realistic if a proactive approach to determine the readiness of the home network for FDX is developed.

#### 4.3.5. Development Progress


Full Duplex DOCSIS requirements began as an appendix to the DOCSIS 3.1 specification and was introduced as a CableLabs project in mid-2016. Recall, it uses the same OFDM and OFDMA-based physical layer for downstream and upstream communication, with the extension of the upstream band to include 108-684 MHz. The Engineering Change Release for this update was in late 2017, and the “cleaned up” spec released in early 2018. Comcast embarked on a node RFP for a DAA-based FDX node and RPD late in 2018 and launched the FDX development program mid-2019. All aspects of development above are underway. In the case of the CMTS, as noted previously, all of our DAA deployments are done with a virtual CMTS (vCMTS), and all FDX will be done via DAA. Thereby, development of FDX features such as MAC management messages, sounding, scheduling, RBA switching, etc., are features under development on the vCMTS.

Current expectations are that 10G FDX end-to-end integration will be taking place in labs across key technology partners, CableLabs, and Comcast in 2021. Operationalizing FDX via tools, training, processes, and integration of services with existing back-office system will also be taking place in order to ready FDX for scalable roll-out. There will be more news to come in 2021 as the integration comes

together and we look towards early sites to prove out 10G FDX technology and multi-gigabit symmetrical coaxial last miles. Stay tuned for Gigamania!

#### 4.4. FDX Spectrum Definition

As would be expected, a significant part of the development of the FDX specification was defining the spectrum plan. Lifting the upper spectral edge of the tried-and-true 5-42 MHz upstream/reverse signal path was no small task. The detailed outcome, published in [17], is the 108-684 MHz FDX band – so-called regardless of whether FDX channels occupy the entire band. Figure 21 shows the band plan definition as written into the DOCSIS 4.0 PHY specification [17].



**Figure 21 - DOCSIS 4.0 FDX Band Plan [17]**

The MUST capabilities as articulated in the DOCSIS 4.0 specification include:

- Allocated FDX spectrum starts at 108 MHz
- 96 MHz of FDX Bandwidth: Located in a single FDX Sub-band from 108 MHz to 204 MHz
- 192 MHz of FDX Bandwidth: Located in two FDX Sub-bands from 108 MHz to 300 MHz, consisting of two 96 MHz downstream channels and two 96 MHz upstream channels
- 288 MHz of FDX Bandwidth: Located in three FDX Sub-bands from 108 MHz to 396 MHz, consisting of three 96 MHz downstream channels and three 96 MHz upstream channels
- 384 MHz of FDX Bandwidth: Located in two 192 MHz FDX Sub-bands from 108 MHz to 492 MHz, each consisting of a single 192 MHz downstream block and two 96 MHz upstream blocks
- 576 MHz of FDX Bandwidth: Located in three 192 MHz FDX Sub-bands from 108 MHz to 684 MHz, each consisting of a single 192 MHz downstream block and two 96 MHz upstream blocks
- The FDX downstream channel and upstream channels sharing the same sub-band must use the same subcarrier spacing and cyclic prefix length. The subcarrier spacing and cyclic prefix on different sub-bands are allowed to be different.

Now, because DOCSIS 3.1 enables a High Split upstream capability defined as up to 204 MHz, and because 1 Gbps is the first step to multi-gig symmetric services, the DOCSIS 4.0 FDX requirements also

accommodate the high split as a starting point. With FDX spectrum beginning at 108 MHz, high-split spectrum will overlap with FDX in the first 96 MHz block, between 108 MHz – 204 MHz.

Left to their own operation, a High Split upstream transmission could interfere with an FDX downstream. This is avoided by ensuring that existing DOCSIS 3.1 modems participate in the sounding and scheduling of the FDX band, although they will not be able to use the FDX downstream band. This is accomplished via the FDX-L software upgrade previously described.

The specification also calls out what are mostly intuitive rules for the use of non-FDX spectrum.

#### 4.4.1. Operator Spectrum Management

Figure 22 shows a typical view of the spectrum for an operator that has launched FDX. The specification basis is a 5-85 MHz Mid-Split upstream and a downstream spectrum to 1218 MHz. When activated, more upstream can begin at 108 MHz by adding 96 MHz at a time. The 85 MHz upstream bandwidth, under today's CAGRs, typical node sizes, and an average utilized bandwidth at peak-busy-hour (PBH), provides an upstream lifespan for many years (5-10 yrs.) when coupled with node splits or N+0. Thus, FDX spectrum is likely to be added to allow new product speeds above the approximate 300 Mbps limit of the 85 MHz band alone.

In Figure 22, the additional FDX bands above 300 MHz can be added as product demands deem them necessary, presuming operators have managed the rest of their DOCSIS 3.0 and video spectrum accordingly. Note that the FDX upstream spectrum can ONLY be spectrum allocated to where DOCSIS 3.1 downstream activities exist, and where that is so, the DOCSIS 3.1 device should be participating as an FDX-L modem.

Lastly, we note that product speeds over an 8-10 year time frame may move from 1 Gbps symmetric to 2 Gbps/2 Gbps, 3 Gbps /3 Gbps, etc. A 3 Gbps /3 Gbps service will extend the FDX band to the 492 MHz block edge (give or take and based on bps/Hz efficiency). This upstream allocation provides the spectrum necessary for the very rare bursts of full peak speeds. However, that spectrum is also available for use by the DOCSIS 3.1 downstream, between these rare bursts. By contrast, in an FDD system such as Extended Spectrum DOCSIS, there is still a need to allocate 492 MHz of spectrum for a 3 Gbps service. However, in this case it is a fixed upstream, set aside to be used only for the very rare upstream bursts. The vast majority of time this prime portion of coaxial spectrum sits idle. To compensate for idle spectrum no longer available for downstream, the entire plant needs to be rebuilt to extend the spectrum, including every active, tap and all in-home splitters, and in so doing also pay the penalty of another lost 100 MHz due to guard band above 492 MHz.

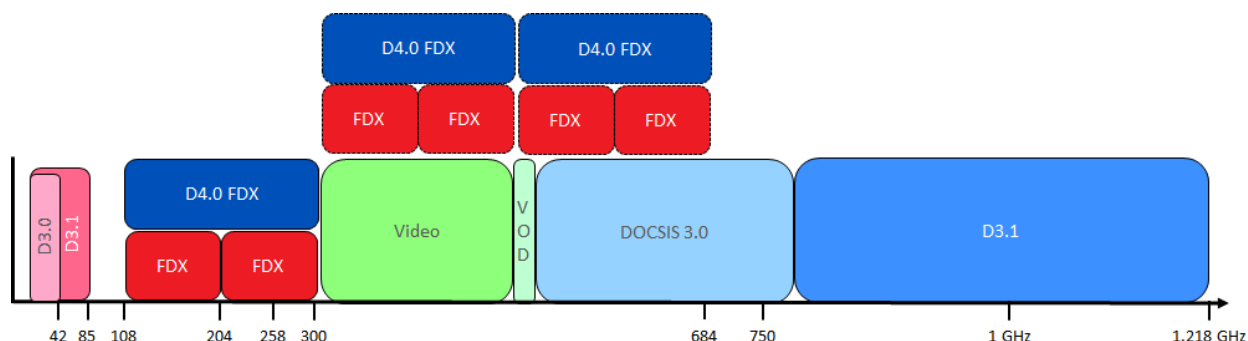



Figure 22 - "Classic" View of Mid-Split Based FDX Implementation

Figure 23 and Figure 24 show some of the practical scenarios to be accounted for as any FDX roll-out begins. In Figure 23, for example, High Split modems may be in the field alongside sub split modems and FDX modems. Often it is best to visualize the spectrum “as the node sees it” and “as the CM sees it,” because the introduction of FDX and the mixing of CM splits on the plant tell a complete story. In Figure 23, we can see that as a node goes from High Split to FDX mode, the spectrum between 85-108 MHz is vacated, removing capacity from High Split only devices. This could matter to devices offering 1G upstream service, since High Split does not provide much margin above 1 Gbps of total capacity. It may, for example, force these CMs to move to an all-OFDMA upstream configuration if they are not already using that, or force the Time and Frequency Division Multiple Access (TaFDM) mode, whereby SC-QAM and OFDMA exist in the same spectrum in different time slots. Lastly, it could mean new FDX CMs for the 1G upstream users, since they are the more likely customers to require multi-gig performance sooner than the rest, and only FDX can go above 1Gbps.

The other item to identify in Figure 23 is that DOCSIS 4.0 downstream is shown in the 204 MHz to 300 MHz range, and a DOCSIS 3.1 downstream above 1 GHz. The line pointing at each is to identify this OFDM block can be in either location. It is to point out that spectrum migration is a living plan, striking a balance between migration complexity and available empty bands to define the sequence of events. For example, if there is no spectrum used above 1 GHz, then adding more DOCSIS 4.0 (just DOCSIS 3.1 at that point) above 1 GHz can defer some of the complexities of overlapping FDX spectrum. Such a determination does not need to be made now, but as more DOCSIS 3.1 is demanded, the overall spectrum plan should be visited regularly against feature roadmap and device penetrations in selection of the optimal path forward.




**Figure 23 - DAA Node in FDX Mode with Mid-Split, High Split and FDX CMs – Mid-Split based Duplex**

Lastly, consider Figure 24, which represents an FDX concept where the node in FDX mode has a High Split duplex instead of a Mid-Split duplex. New FDX spectrum is added “on top” of the High Split spectrum, as shown. There may be areas built only as High Split, such that an FDX upgrade that places 4+ FDX upstream blocks above the High Split band would enable current devices to operate as they already do, while allowing customers seeking >1G speeds upstream to obtain an FDX CM, and access the FDX channels beginning at 258 MHz.

As is common when engineers are tasked with defining what is often the far future when writing technical specifications, the specification’s wording is ambiguous for anyone seeking to “comply to the letter” of its definition. However, this may be an implementable option and is currently being evaluated as a viable deployment model. As pioneering work on FDX integrations get underway, it is likely that many specification interpretations and meanings will need to be worked out among collaborating partners. This is a common component of technological advancements.






**Figure 24 - FDX Node Concept as a High Split Augmentation**

#### 4.5. More DOCSIS 4.0

The 10G initiative is based on understanding the capabilities of the HFC network, and the emergence of key technologies and applications of the future. However, “10G” does not explicitly describe a “how-to.” This is purposeful, because differences exist among MSO starting points and roadmaps. Different tools in the toolbox inevitably occur over time. A good example of this was different approaches taken to deal with the growth of HD television – analog spectrum recapture using Digital Terminal Adaptors, switched digital video (SDV), and/or 1 GHz expansion all were used.

Different approaches are also available for 10G. The DOCSIS 3.1 specification calls out optional use of spectrum to 1.794 GHz (which goes conversationally by 1.8 GHz). However, at the time (2012), the system engineering was deferred at that time, as this was not a 2012 priority. This has now been completed as part of the DOCSIS 4.0 specification and is referred to as “Extended Spectrum DOCSIS” or Frequency Division Duplex DOCSIS (FDD).


FDD goes beyond the original DOCSIS 3.1 upstream also, defining an upper spectral edge for upstream connectivity beyond 204 MHz, much as FDX does. The FDX upstream grid and current 96 MHz block OFDMA definition enables synergy between FDX and FDD. This synergy is shown in Figure 25.



**Figure 25 - DOCSIS 4.0 Full Duplex and Extended Spectrum Increase DS OFDM and US OFDMA by Different Means**

As shown in Figure 25, a system design that supports both FDX functions and a 1.8 GHz downstream upper edge can have a common number of downstream OFDM and upstream OFDMA blocks. This is beneficial for silicon solution vendors when developing signal processing chains and sizing chip resources. However, they are allocated differently on the RF coaxial plant.

In a potential “DOCSIS 4.1” future, we can envision the two DOCSIS 4.0 technologies of FDX and FDD as complementary. The combined power of FDX and FDD provides yet another avenue for capacity expansion on the network. This is shown in Figure 26.



**Figure 26 - DOCSIS 4.0 Full Duplex and Extended Spectrum are Complementary in Nature**

## **4.6. FDX Amplifiers**

### **4.6.1. Echo Cancellation-Based Amplifiers**

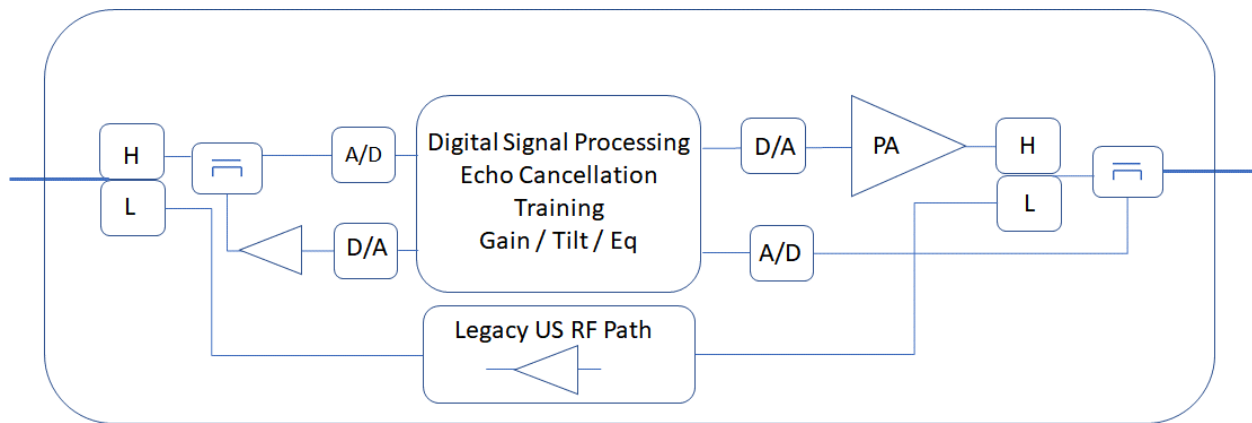
The FDX specification was written under the assumption of an N+0 network. This simplifies the introduction of FDX by getting fixed diplexed amplifiers out of the way, and it aligns with an architecture path some operators are taking. However, it is not technically limited to an amplifier-free plant. The EC technology developed can apply at any point in the network to manage overlapping spectrum, and this can include amplifiers. Of course, these are not traditional amplifiers, but a new class of device that includes this essential signal processing.

The system engineering aspects of EC-based FDX amplifiers includes:

- 1) How good does the EC need to be?
- 2) Amplifiers can have an “amplifying” effect on IGs, depending on the HFC design. How important is this?
- 3) How does the amplifier specification change with respect to RF performance? Are there cascade limitations?

An EC-based amplifier concept is shown in Figure 27. The nature of overlapping spectrum and gain in both directions creates a full-circle loop gain path, by design. There is a concern even in normal HFC amplifiers that indirect loop gains occurring in certain frequency bands could cause oscillation. The EC must, as a minimum, suppress this gain loop, meaning net gain around the path must be < 0 dB.






**Figure 27 - An Echo Cancellation-Based Amplifier to Support FDX Beyond N+0**

Echo-cancellation operates similarly to that of the node. A sample of the upstream signal is taken so that an opposite phase, equal magnitude version can be added in front of the downstream amplifier in Figure 27. Similarly, the signal from the downstream amplifier is sampled and an anti-downstream version added at the input to the upstream amplifier. As in the node, the EC adjusts for changes on the echo response.

Analysis of the EC performance necessary was calculated for two scenarios, based on common HFC amplifier RF level and noise parameters, and shown in Figure 28 [16]. The left figure shows the use of overlapping spectrum for the entire FDX band without constraints. In this case, the objective is to maintain high enough SNR (minimal residual EC degradation) to support the OFDM and OFDMA formats defined in DOCSIS 3.1, with the assumption being up to 2048-QAM in the downstream and 1024-QAM in the upstream, and assuming a reflection magnitude of -15 dB exists at each port. The EC performance requirement, shown along the X-axis, is about 75 dB before residual echo interference begins to impact the SNR required for 1024-QAM. Using the dB relationships between formats (3 dB per bit), these charts can be extrapolated in either direction (for more or less bandwidth efficiency) to arrive at EC requirements for other reflection or SNR assumptions under existing noise assumptions. For example, for a less-efficient 256-QAM upstream, which has a 6 dB lower SNR threshold, the EC requirement reduces to about 65 dB. These are challenging requirements, particularly for a single-stage echo-cancellation design.

For the second case, the amplifier is not required to support the simultaneous use of shared spectrum. In this case, the goal of the EC is to assure that the amplifiers are able to operate sufficiently in a linear mode, thus not degrading noise and distortion. This case, shown on the right, indicates that 25-30 dB is sufficient, which is a significant relaxation compared to the first case.



**Figure 28 - Worst Case SNR vs. Echo Cancellation [16]**

What this analysis points out is that FDX amplifiers are likely to be a tradeoff of complexity, cost, efficiency, HFC design, and also SW complexities associated with the additional EC to be trained and to manage. As well, possible refinement of IG/TG algorithms associated with RBA assignments to include the impact of FDX amplifiers.

#### **4.6.2. RF Isolation-Based Amplifiers (No EC)**

Because HFC has always been an FDD system, diplex filters have always been used. Their purpose is to isolate downstream and upstream signal paths from each other. Diplex splits have evolved over time from 42/54 MHz to 65/85 MHz, 85/102 MHz, and 204/258 MHz. In the future, with Extended Spectrum DOCSIS, higher splits are expected.


The downside of diplexers are that when the split needs to be changed, it is labor intensive to change all of them (2 per amplifier) or swap active components. In addition, valuable spectrum – 20-25% of the band edge – is lost in the transition band to make filters of practical cost, size, and repeatability. The amount of MHz lost to diplexers increases with absolute band edge for the split. For a 492 MHz FDD split, it's about 100 MHz of lost spectrum!

An ideal solution would be an amplifier with NO diplex filters. This would make it feasible to build a transparent, flexible network with zero-touch for upstream/downstream spectrum re-allocation, and no guard bands.

There are two significant challenges to a diplexer-less amplifier:

1. How do you create an amplifier with loop gain below 1 to prevent oscillation?
2. How do you create a system able to work with real world return losses?

Consider Figure 29, which looks at the input port of a diplexer-less amplifier, where in their place are specialized multi-port directional couplers (MC x).



**Figure 29 - Evaluating the Signal Relationship of a Diplexer-Less Amplifier (courtesy Technetix)**

We define the following (all in dBmV):

A = input signal

B = (input signal A) – (insertion loss of the coupler MC1)

C = B + (gain Amp1)

D = C – (insertion loss of the coupler MC2)

E = (leakage of C in MC2)

F = E + (Gain Amp2)

G = F – (insertion loss of the coupler MC1)

H = (leakage of F in MC1)

Then, looking at the amplifier output to the downstream (point D):

$D = (\text{input A}) - (\text{Isolation MC1}) + (\text{Gain Amp1}) - (\text{Isolation MC2.})$

Thus, the amplifier works only if:

(1)  $H < B$  :


$$(\text{Gain of Amp1}) + (\text{Gain of Amp2}) - (\text{isolation MC2}) - (\text{isolation MC 1}) < 1$$

(2)  $RL = G < A - 20\text{dB}$  :

$$(\text{Gain of Amp1}) - (\text{isolation of MC2}) + (\text{Gain of Amp2}) - (2 \times \text{IL MC 1}) < -20\text{dB}$$

An amplifier that can achieve both (1) and (2), can be a transparent, diplexer-less amplifier. By combining an amplifier with these two operating principles, along with a 54 MHz high-pass filter (HPF)

and a 684 MHz lowpass filter (LPF) – the edge of FDX or Extended Spectrum DOCSIS upstream – the frequency response behavior shown in Figure 30 is obtained.



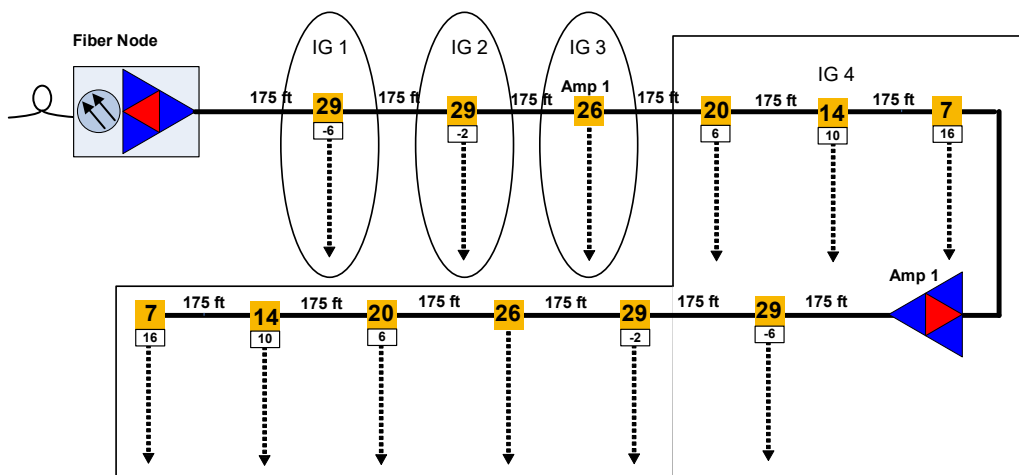
**Figure 30 - Forward and Reverse Freq Response Characteristic of a Diplexer-less Amplifier (Courtesy Technetix)**

In Figure 30, the purple area shows the region for which there is gain simultaneously in both the upstream and the downstream signal paths, without diplex filters in place. This creates a uniquely transparent and flexible network between the node and the customer premise, either by passing FDX signals seamlessly or by supporting any desired split that does not exceed 684 MHz in an Extended Spectrum FDD system.

From an FDX system engineering standpoint, the diplexer-less amplifier results in an extension of the echo cancellation capability in the time domain. For example, if a 5-tap string becomes a 10-tap string, reflections occur from further away points from the EC circuit. Typically, long-distance reflections also attenuate in magnitude with time. This will not be the case with gain in the return path that will forward along a reflection that is on the “south” side of the amplifier. This makes the job of the EC more challenging.

#### **4.7. Traffic Engineering for FDX Services**

Consider the case of an N+1 system based on FDX and shown in Figure 31. When an amplifier is included, using current tap RF performance, one of the effects is the expansion of Interference Group 4 (IG4) to the “south,” or home-facing, side of the amplifier. These users become part of the last IG of the tap string before the amplifier. This is because of the limited drop-to-output isolation characteristics of today’s taps.




**Figure 31 - Potential Interference Group Expansion due to Amplifier on an FDX Network**

Taps with port-to-output isolation like those with higher tap values could eliminate this phenomenon. We have developed tap specifications with FDX requirements in mind.

Nonetheless, since current taps will be in the field for years, leading to enlarged IGs, traffic engineering analysis was performed to understand, quantifiably, the sensitivity of FDX performance to the size of an IG, the total spectrum, the service group size, speed tiers, etc. The empirically-based modeling is founded on real network traffic distributions gathered on production CommScope CMTS platforms over many years. Figure 32 shows the components of the traffic useful for helping to understand the problem and visualize the results [4].

The traffic can be broken into three components [4]:

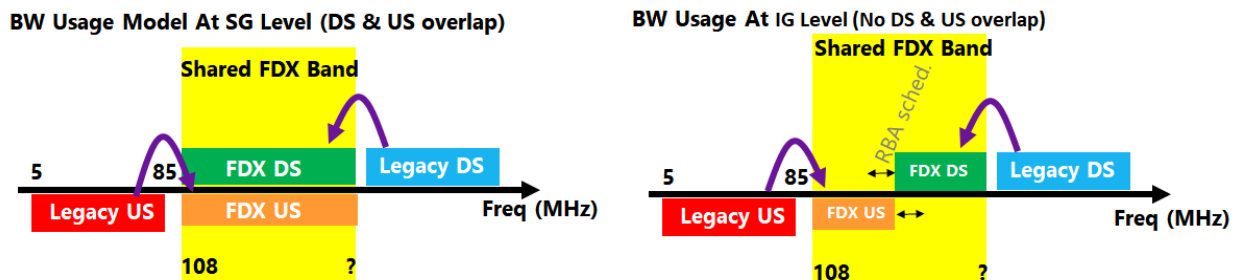
- 1) Average Service Group (SG) Utilization = Number of subscribers in the service group x average utilization per subscriber at PBH.
- 2) Ripple = variations around the average utilization because traffic is bursty and not a constant stream equal to the average. Though statistically imprecise, the ripple is somewhat like standard deviation. However, it also encompasses a component tied to the probability of bursting to peak speed.
- 3) Max Burst = Just that, the moments when the traffic burst is 100% max'd out at the peak speed offering of the high speed data (HSD) product in this service group



**Figure 32 - Breakdown of HSD Traffic Components [4]**


HSD product speeds directly translate into a minimum spectrum allocation needed to guarantee a maximum speed burst. Operators commonly allocate peak spectrum and an empirically derived additional capacity that accommodates the three components of Figure 32. On real networks, rules have been developed that link service group size, speeds, and percent capacity utilization to define the total capacity requirement. Then, with an awareness of device penetrations, this can be translated to DOCSIS 3.0 and DOCSIS 3.1 spectrum requirements. These rules were re-calibrated again with the introduction of DOCSIS 3.1 and 1 Gbps downstream speeds.

All of the above analysis and empirical rule making apply to FDX – it is still about downstream and upstream spectrum requirements to support certain product speeds. What changes for FDX is that the FDX band is allocated for both downstream and upstream. The infrequent bursts of peak speeds can therefore be called upon to service the downstream or the upstream, and they are not completely independent because of the IG phenomenon. Although from the node perspective there is full duplex spectrum operation, from the CM perspective the introduction of IGs places constraints on who can simultaneously access the spectrum. Figure 33 depicts this observation.



**Figure 33 - Managing FDX Bandwidth to Guarantee Peak Speed Bursts [4]**

Shared spectrum does not automatically mean peak bursts within an IG will compete for FDX spectrum. It depends on how much FDX spectrum is allocated, the distribution of peak speeds, the maximum speed offering, and of course the size of the IG and the number of subscribers in general. However, the FDX band carries around 5 Gbps. So, as speed tiers increase, at some point, from the perspective of an IG, there may not be enough FDX spectrum to avoid downstream and upstream burst demand exceeding available FDX band, as depicted in Figure 34. Naturally, it is much more likely to be needed for the upstream, since there is a substantially larger downstream allocation for DOCSIS 3.1.



**Figure 34 - Representation of Colliding Simultaneous DS and US Bursts of Shared Spectrum [4]**

Some context may be helpful at this point. Development of the IG/TG concept was formulated during the early specification discussions of FDX. DOCSIS bandwidth has always been shared in the downstream and in the upstream with great efficiency by using relatively large service groups. For FDX development, instinctively there was a thought that having an IG as small as reasonably possible was a sound goal and would not significantly change the capacity utilization dynamics. A tap being its own IG, or several taps being part of the same IG, is what occurs using typical RF performance. There were not major concerns about the impacts of such small IGs; there was more concern about the process of sounding the network to define them.

The concern for IG/TG size increased when the concept of the FDX amplifier came into the picture to support FDX in non-N+0 systems. It was during system analysis with amplifiers that the IG elongation phenomenon shown in Figure 31 was uncovered. This led to questions:


- 1) How large can an IG be before there is an impact to the customer experience?
- 2) What service speed / IG size / spectrum rules exist when downstream and upstream traffic engineering become co-mingled in FDX?

We set out to find answers to these questions, working closely with experts with sophisticated HSD modeling tools at CommScope.

Figure 35 shows a model assessing the subscriber count, spectrum, QAM efficiency trade space for a set of input parameters projected for 2028 that includes a post-Covid average utilization baseline, assumed capacity CAGRs of 35% downstream and 30% upstream, and speed tiers up to 4 Gbps/4 Gbps. In these models every subscriber is given this peak service tier. Although the highest speed tiers are, in fact, typically the lowest penetrated, this obviously makes for a conservative analysis.

What can be concluded from this analysis and Figure 35 is that the subscriber group size (IG size) sharing FDX spectrum from 108-684 MHz can be as large as 64 (red circle) and stay within 1200 MHz of total spectrum, or even within 1 GHz for a 9 bps/Hz net downstream throughput average efficiency. The upstream, even for the 4 Gbps/4 Gbps (green) case, stays within the 576 MHz of maximum FDX band allocation, even for an IG size above 128.

What these studies reveal is that FDX bandwidth can be used extremely efficiently, shared across larger IGs than perhaps initially envisioned, and the statistical nature of peak bursts is better understood. What we have learned is these peak bursts are very infrequent per user, and the collision of bursts from concurrent peaking users is more infrequent still. With an 85 MHz legacy upstream, the upstream bandwidth is sufficient for PBH average utilization for a subscriber size of up to about 150-200 subscribers, shown in Figure 25 as approximately where the red arrow points to the FDX bandwidth breach. Thus, allowing the FDX upstream allocation to be a bandwidth reservoir for peaking users while still serving downstream capacity is an extremely efficient use of precious HFC spectrum, compared to setting aside hundreds of MHz of the highest quality HFC to, mostly idle.



**Figure 35 - Total RF Spectrum Required vs Subs Sharing an IG for 3 Gbps/3 Gbps, 4 Gbps /2 Gbps, and 4 Gbps /4 Gbps [4]**

## 4.8. Operationalizing FDX

### 4.8.1. Echo Cancellation and Diagnostics - Live!

As a starting point for understanding EC math, consider first a non-FDX signal. Let  $y(n)$  represent a noiseless signal presented to a receiver, which is a summation of the received signal, and its attenuated and delayed signal copies of delay,  $D$ .

$$y(n) = \sum_{d=0}^D b_d x(n-d) \quad (\text{non-FDX Signals})$$

The signal copies come from plant defects including micro reflections, and other frequency response effects like filter roll-off and group delay variation (GDV). Each delayed copy will be attenuated and phase-shifted by a complex coefficient,  $b_d$ .

Non-FDX signals can be made quite robust, thanks to equalization technology available in DOCSIS today, enabling an efficient means of adapting equalizer coefficients that compensate for  $b_d$ , and deliver  $y(n) \approx x(n)$ .

By contrast, an FDX system will have both downstream and upstream signal components. At the node receiver, since signals overlap with no filters to separate them, the received signal  $y(n)$  is:




$$y(n) = y_{ds}(n) + y_{us}(n) \quad (\text{FDX Signals})$$

The FDX echo canceller must estimate both the downstream signal,  $b_0 x_{ds}(n)$  and its echo profile,  $\sum_{d=1}^D b_d x_{ds}(n-d)$  to minimize the downstream contribution at the upstream receiver such that  $y(n) \approx y_{us}(n)$ , where:

$$y_{ds}(n) = \sum_{d=0}^D b_d x_{ds}(n-d)$$

In addition to echo-cancelled MER and the FEC, the coefficient  $b_d$  can assist the MSO in understanding the impairments at the FDX receiver at any given time, similar to the familiar linear distortion impairment analysis that can be obtained from upstream equalization [15]. EC analysis can lead to a wealth of PNM knowledge previously unavailable because the non-FDX or traditional node receiver echo profile view was blocked by a diplex filter.


A sample echo profile is shown in Figure 36 in the brown trace labeled “Vendor EC”, representing a snapshot of what an echo canceller would need to compensate for, when connected to a cable plant where the RF taps are spaced between varying lengths of cable (90-132 ft). If a time domain reflectometer (TDR) were to connect to the plant at the same location as the node, it would produce a similar echo profile, shown in the bottom blue trace labeled “VNA TDR.”



**Figure 36 - Echo Profile from New London, CT FDX Field Trial**

In Figure 37, the shaded block illustrates the coupling of a vector signal analyzer (VSA) and DOCSIS cable load generator (CLGD) via a resistive (low isolation) splitter at the same location of the node used in Figure 36. All of the remaining traces, labeled “VSA-156 MHz” through “VSA-636 MHz” are fully

adapted VSA equalization coefficients,  $b_d$ , at all of the FDX center frequencies associated with a 96 MHz OFDM block transmitted by a CLGD and received by a VSA. All VSA traces mostly overlap with one another except for minor differences at each of the tap locations, where it is clear the overlays are visible.



**Figure 37 - CLGD/VSA Echo Profile Coefficient Capture**

Diurnal and seasonal trends typically result in small and nearly continuous response variations over time. This includes changes in either temperature, or other weather-related events including wind loading. EC response data can tell us the same thing we'd learn by connecting a TDR at the node location. They can correlate to a plant map's precise location of every component, including taps, passives, splices, and cables. Each spike in the echo profile of Figure 36 corresponds to each tap within the feeder leg shown in Figure 38. The distance to the fault value of each spike directly relates to the length of cable between the node and each of the taps, specifically by a factor of 2. So, 90 ft between the node and the 2nd tap – a 2-port, 29 dB tap in Figure 38 -- is represented as approximately 180 ft from the fault measured in Figure 36.




**Figure 38 - CATV Plant Map Corresponding to Figure 36 Echo Profile**

It is expected that it will be reasonably easy for EC technology to adapt to these static or slow-moving echo profiles. Transient behaviors, on the other hand, can represent a more stressful scenario for EC. Plant maintenance activities alone can result in rapid and potentially discontinuous response variation over time. Some common plant maintenance activities include:

- a) Tap face plate removal and insertion
- b) Interference testing
- c) Seizure screw checks (loosening and tightening)

Tap port termination quality was evaluated as well and showed negligible impacts to the echo profile.

Figure 39 shows how the reflections behave on a millisecond scale, as a technician removes or installs a tap face plate. When the faceplate is removed, an internal continuity switch occurs, resulting in a spike representing a nearly 30 dB change. In these cases, the EC will need some time to adapt to the abrupt changes, and both MER and FEC readings will also reflect this transition. However, MER and FEC will not be able to identify the precise location of where this change occurred. This is precisely where EC may provide MSOs with new value.



**Figure 39 - Single Coefficient Variation from Faceplate Removal and Insertion**

A potentially new operations dashboard may abstract away many of the gory details associated with analyzing EC echo profiles, perhaps only raising awareness when the EC has moved into a more aggressive mode of adaptation, due to a breach of a discontinuity threshold. This mode mostly provides confirmation of ongoing plant maintenance activities, but it can also assist MSOs in recognizing nearly equivalent changes occurring outside of planned plant maintenance activities. Those alerts, combined with precise location of the fault, could represent an unprecedented level of Proactive Network Maintenance.

#### **4.8.2. Proactive Network Maintenance**

Speaking of Proactive Network Maintenance (PNM), it has been an integral part of the DOCSIS specification since DOCSIS 3.0. For nearly 15 years, cable operators have been learning how to harness the vast capabilities available within the PNM suite of tools. With the addition of Machine Learning and Artificial Intelligence (ML/AI), operators are gaining unprecedented visibility into the echo profiles of their coaxial distribution networks. These tools enable operators to remotely evaluate the spectral performance at nearly every location of plant that has a cable modem installed.

The (partial) list below summarizes why operators appreciate what PNM brings to the network reliability and positive customer experience fronts. The first four help to optimize the channel performance in the face of network variations that are different from one end-to-end connection to another. The bottom four are powerful tools for monitoring, identifying, troubleshooting, and remediating problems found.

- Upstream adaptive pre-equalization
- Channel estimate coefficients
- Downstream equalization coefficients

- RxMER per-subcarrier analysis
- Full Band Capture (FBC) spectrum analysis
- Spectral Impairment Detection (SID)
- Triggered upstream spectrum capture
- Active and quiet probes

What do these PNM capabilities all have in common? Most of them can play at least some role in directly or indirectly measuring return loss and reflected energy within the coaxial plant. These tools can be used together with system design maps to inform echo cancelers in FDX deployments.

Typically, the PNM tools are used to detect, analyze and troubleshoot problems within their respective RF spectrum and direction. For example, in the case of upstream adaptive pre-equalization, DOCSIS station maintenance requests are used to continuously monitor and adapt the upstream channel for micro-reflections (echoes). Our PNM systems provide a high-resolution analysis of DOCSIS channels in use. In the case of a typical 24-tap SC-QAM pre-equalizer, there are resolution limits to the equalizer. For example, a 6.4 MHz wide SC-QAM channel will have a symbol rate of 5.12 Mbps, resulting in approximately 213.33 kHz resolution bandwidth (RBW). However, using DOCSIS 3.1 OFDMA upstream pre-equalizer coefficients, RBW will be as high as 25 kHz, and typically 50 kHz.


Collecting, analyzing and grouping the pre-equalizer response signatures results in an exceptional system for locating small pockets of network impedance problems that cause low frequency echoes. Figure 40 illustrates the effectiveness and accuracy of this technique. While the current DOCSIS upstream spectrum operates at lower frequencies than FDX, this will undoubtedly play an important role in troubleshooting and maintaining our physical plant. Having the echo response at these low frequencies enables operators to pinpoint virtually every impedance mismatch on the outside plant. It's also important to note that pre-equalization is available on all DOCSIS compliant devices since version 1.1.



**Figure 40 - Upstream Adaptive Pre-Equalization Analysis; Echo Response Groups, Mapped**

Arguably the most useful PNM capability for troubleshooting RF performance is downstream Full Band Capture (FBC). The previously described technique of the pre-equalizer response grouping has been extended to downstream spectrum analysis. Figure 41 illustrates a common example of a standing wave,

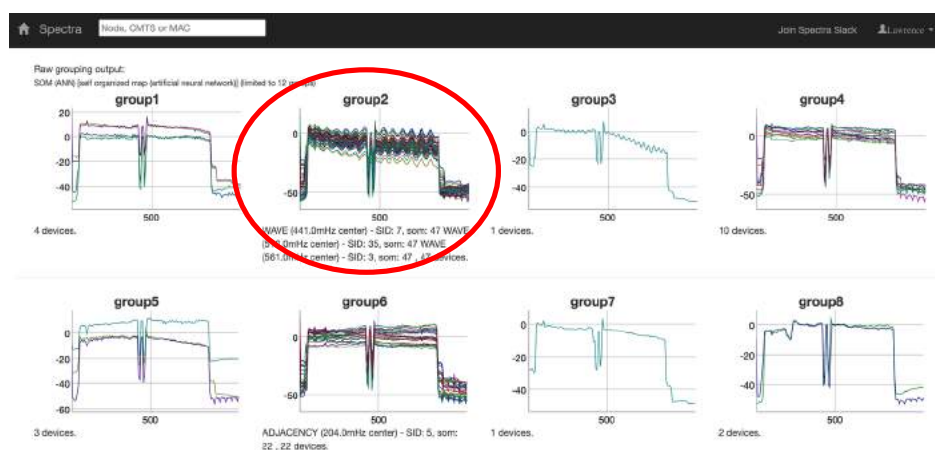
Spectra



### Figure 41 - Downstream Full Band Capture, Standing Wave

26 MHz to 1026 MHz at 117.18 kHz RBW

Similar techniques of signature matching based on Artificial Intelligence (AI) prove very effective at grouping common response issues. Figure 42 shows that the standing wave signature is localized to a group of 47 similar frequency responses. These 47 cable modems share a common 1842 foot .875 P3 trunk cable with a Velocity of Propagation (VoP) factor of 0.87. Having a peak-to-peak frequency separation of 47 MHz, the echo distance can be calculated using the speed of light in feet-per-second:  $(983,571,056 / 2) * 0.87 / 47 = 9$  feet. In the case of this wave, the higher attenuation at lower frequency indicates that the primary reflector is near the launch end of the cable and the secondary reflector is most likely at the output of the node. In this example, the echo profile of this faulty cable can now be calculated in both directions, at any frequency within its operating range. Relying on well-known principals of antenna reciprocity, the important scattering parameters (S-parameters) can all be calculated for this network segment.



**Figure 42 - Artificial Neural Network, Response Grouping; Self-Organized Map (SOM)**

In summary, PNM tools have been around for a long time and have become extremely sophisticated at pinpointing plant issues, and typically before a customer would notice an impact. FDX relies by its very



nature on overlapping spectrum and knowledge of the plant echo characteristic over time. The advances in PNM align perfectly for the introduction of FDX, and the introduction of FDX will further empower PNM through the availability of EC coefficients, as well as the sounding and isolation information developed for FDX operation.

#### **4.8.3. DOCSIS 3.0 and DOCSIS 3.1 Participation in FDX Sounding**

Channel “sounding” is an essential tool for FDX. Cable modems transmitting at high levels in higher frequency spectrum will create new opportunities for co-channel interference (CCI) and adjacent channel interference (ACI) where isolation is insufficient. The ability to coordinate a sounding routine across all modems will be used to assign interference groups (IG). Similar to the PMA in DOCSIS 3.1, there will be opportunities to externalize the sounding routines, allowing for more context and better-informed IG modeling.

For example, a typical FDX deployment will co-exist with a majority population of DOCSIS 3.0 and DOCSIS 3.1 modems. In the case where a customer would like symmetrical multi-gigabit speeds, requiring an upgrade from DOCSIS 3.1 to DOCSIS 4.0 FDX, they will receive a new cable modem/gateway to replace the existing device. The condition of that customer’s network and its upstream impact on adjacent customer communications will not be known until after the new modem is connected and sounding has occurred.

This scenario can be addressed by extending the sounding routines and adapting them to use the external PNM infrastructure. Having full band DOCSIS 3.0 and 3.1 spectrum analysis capabilities within the cable modems, the existing equipment can participate in the sounding routine. There are limits to using the log-magnitude data reported by the full band capture, such as lacking phase information. However, in early years of FDX availability, this type of analysis will provide useful information where FDX sounding clients may be sparse or non-existent.

#### **4.8.4. “Inverted” Plant of FDX Band**

Along with FDX, and like the High Split, comes a new paradigm of high-power transmission in higher frequency spectrum from the CM. In addition to CCI and ACI, there is another potential problem from these “inverted” levels in the plant. Due to dissimilar carbon content in the metallic cable components, and in the presence of electrical current, ionic migration is constantly occurring. In other words, corrosion is commonplace, and is especially exacerbated in the presence of water.

Corrosion in the form of a molecular junction is known to behave similarly to a diode and is often referred to as a corrosion diode. These corrosion diodes can function in unpredictable ways, often having resistive, capacitive and cascaded diode properties. One of the most common side-effects is a non-linear mixing effect that causes 2nd and 3rd order RF products. In traditional coaxial networks, this is known as Common Path Distortion (CPD) and typically occurs near the output of amplifiers, where composite RF power is high. In digital cable systems, the downstream channels become down-converted and mixed with the upstream frequencies, which are much lower in RF spectrum. The summation of the 2nd and 3rd order products typically manifests as an elevated noise floor in the upstream spectrum.

It can be anticipated that a significant number of corrosion diodes may exist in our collective drop networks, which tend to be older and less well maintained than the hardline portions of the plant. However, with the relatively low power levels of downstream RF in the drop cables, there is insufficient power to forward-bias any diodes that may exist. In the case of High Split or FDX, there will be higher power in the drop cables, which could begin activating countless dormant corrosion diodes. The resulting

2nd and 3rd order products could have significant impacts on performance. This phenomenon is referred to as Passive Intermodulation (PIM).

Fortunately, corrosion diodes tend to create impedance mismatches that can be detected and located using some of the previously mentioned PNM analysis techniques. The SCTE Network Operations Subcommittee for PNM has an ongoing workstream to evaluate the effectiveness of these tools for locating potential PIM sources.


#### **4.8.5. Field Tools**

The fundamentals of cable haven't changed much since the advent of forward and return path RF, separated and protected by diplex filters. Our protocols have been designed to provide robust and reliable service, but human troubleshooting and maintenance are still required at times. FDX introduces a new level of complexity and data analytics that will become increasingly reliant on automations and algorithms to augment technicians existing tools for troubleshooting and repair. Many of the tried-and-true techniques for isolating service problems will have limited use in a world of bi-directional RF, and new processes will be developed and refined as FDX rolls out

## **5. Bringing Low Latency DOCSIS to Life**

Without latency and jitter improvements, just delivering faster speeds will not guarantee the best customer experience. Low latency is pivotal to residential services such as gaming and video conferencing and are almost always tied into service level agreement (SLA) requirements for commercial services applications. The expectation that low-latency services will become increasingly important also revolves around AR/VR, IoT, and machine-to-machine (M2M) communications. Higher downstream and upstream utilization levels have an impact on latency, jitter and packet loss, that cannot be resolved only by increasing speed. These metrics and measurement techniques, in addition to speed and availability, must be integrated into design and operations and turned into actionable data. Although speed tests have been used as a measure of performance, other test tools tests have not yet been integrated into operations. Recent architecture changes have enabled MSOs to measure latency under load and without load. In this section, we will provide techniques and example performance that can be achieved in today's networks. We will then describe emerging features to deliver lower latency services as part of the 10G roadmap.

Current speed tests first check connection latency. Since the speed tests are performed during relatively idle utilization periods, connection latency measured between the customer router and test server within the MSO's network includes mainly basic DOCSIS media delay, as shown in Figure 43. For idle utilization levels and good network conditions, latency and jitter distributions do not have heavy tails of large values in their histogram.



**Figure 43 - CDF of RTT Measurements Under Low Utilization Levels for a Mix of Different CM Types and Speed Tiers**

However, as shown below in Table 1 (<https://cablela.bs/low-latency-docsis-technology-overview-february-2019>), queuing is the largest source of delay when the home network is congested. Both downstream (CMTS) and upstream (CM) HSD service flow queues are single queues. Low latency services such as gaming may be delayed due to concurrent queue building (QB) applications, such as file uploads and downloads. It's the networking equivalent of "take a number." Furthermore, unmanaged large queues can lead to the "buffer bloating" effect. Finally, bonding group utilization and network conditions may delay packet transmissions, increasing queueing latency.

**Table 1 - Delay Sources in DOCSIS Networks**

| Delay Source           | Upstream            | Downstream           | Total               | Notes                                     |
|------------------------|---------------------|----------------------|---------------------|-------------------------------------------|
| Queuing                | 0 – 100+ms          | 0 – 100+ms           | 0 – 200+ms          | Largely caused by TCP                     |
| Media Acquisition      | 2ms – 8ms           | 0                    | 2ms – 8ms           | Request/Grant process w/ 2ms MAP interval |
| Serialization/Encoding | 0.38 – 2.8ms        | 60µs – 720µs         | 0.44ms – 3.5ms      | Based on Channel Configuration            |
| Propagation            | 8µs - 300µs         | 8µs - 300µs          | 16µs - 600µs        | 8µs/mile one-way based on furthest CM     |
| Switching/Forwarding   | 1-20 µs             | 1-20 µs              | 2-40 µs             | Implementation dependent                  |
| <b>TOTAL</b>           | <b>2.4 – 111+ms</b> | <b>68µs – 101+ms</b> | <b>2.5 – 222+ms</b> |                                           |

Queuing algorithms vary widely among current CMTS and CM models, and DOCSIS versions. Buffer Control in D3.0, and Active Queue Management (AQM) in D3.1, aim to remove buffer bloat. Latency under load (LUL) measurements reflect the maximum latency/jitter a customer may experience under heavy home network utilization. Generally, these tests are done under downstream and upstream load, sequentially. Bi-directional tests are also used to measure the impact of downstream load on the upstream latency and vice versa. The FCC's "Measuring Broadband America" program, as well as third-party ISP rating tools, use LUL as a performance indicator.




If the queue size is not managed, customers may be affected by buffer bloat. An example is shown in Figure 44, which has a large queue setting in the CMTS. Traditional CMTSs have large physical queues. In this example, higher downstream speed tiers have downstream LUL smaller than 500 ms, compared to lower downstream speed tiers that can have ~2 sec maximum downstream LUL (top left) and 1-2 sec mean downstream LUL (bottom left). After optimizing downstream queue parameters, the same network has improved latency without degrading speed test results (top right, max LUL and bottom right, mean LUL).

Note that Figure 44 includes outliers and speed test failures that may be due to test and network issues in the test environment. Figure 44 also shows that although most upstream LUL readings are < 30ms, it may reach >100 ms depending on the upstream speed tier, buffer control, and queue management algorithms in cable modems. For example, CMs with buffer control where a 250ms default target latency is used have upstream LUL of around 250-300ms, much higher than CMs with optimized buffer control or AQM settings.



**Figure 44 - Maximum, Minimum and Mean DS and US LUL with (Left) High Target Latency Values for DS AQM and (Right) Optimized Target Latency Values for DS AQM**

Figure 45 displays the Cumulative Distribution Function (CDF) for round-trip-time (RTT) measurements between the CPE and measurement within the MSO network under real traffic conditions for a given test case. As expected, the latency values are between the idle latency shown in Figure 43 and the latency under load reading shown in Figure 44. The aggregate values show that the 85<sup>th</sup> percentile has RTTs of less than 50 ms. Although high speed tiers help to improve packet delay, the delay is not bounded, and packet delay variation (jitter) due to fluctuations in home network utilization, with bursty queue-building traffic, can cause a degraded experience. These bursts may be short lived and may have little impact on aggregate values, but the spikes may have a significant impact on services like gaming, videoconferencing and AR/VR applications.



**Figure 45 - CDF of Connection Latency Based on TCP Handshake RTT Measurements Under Customer Traffic Load for a Mix of Different CM DOCSIS Versions and Speed Tiers**

Additional DOCSIS 3.1 LLD specifications aim to support bounded latency with three main new functionalities:

- Dual queue coupled AQM with queue protection
- Reduced MAP interval
- Proactive Grant Scheduling

Early results show that dual queue AQM can provide <10ms RTT for 95-99th percentile of packets for latency sensitive services. Proactive grant scheduling (PGS) aims to support ~1 ms RTT the 99th percentile of packets for latency sensitive services.

Deployment of a dual queue AQM approach requires other changes in the current architectures, including low latency service classification and markings integrated into operations. In addition to access networks, home (e.g. Wi-Fi) and core networks need to support these features for end-to-end support.

PGS requires more algorithmic development to optimize efficiency gain vs overhead with real traffic and to understand deployment costs and operational models.

## 6. 10G Security Initiatives

The 10G vision is most often associated with the capacity behind the “10 Gigabit” nomenclature. However, low latency is high among the 10G pillars, driven by gaming applications and more recently with the rapid growth of video conferencing, driven by the pandemic and work-from-home mandates/recommendations. However, low latency has a broad range of residential and commercial applications beyond gaming and Zoom/Teams/pick your most representative video conference app. The explosive growth in M2M and IoT are good examples.

The explosion of IoT devices and data, and other future applications, draw attention to another ever-more-critical area for customers – security. The smart home, security cameras, telemedicine, business data, and

industrial networks present services and application use cases with a high premium on data confidentiality and on the overall security of access and communications. Customers are hyper aware of the need, but at the same time don't want to think about it, and just expect it to be there, 24/7.

To consolidate a common 10G security vision across operators, CableLabs has several parallel programs targeting essential elements that look at both new applications as well as emerging technologies for operators' networks [18].

On behalf of operators, CableLabs is participating in key IoT Security forums in the Internet Engineering Task Force (IETF), National Institute of Standards (NIST), and the Open Connectivity Foundation (OCF), among others. Security of IoT services and devices is a common ground across the device and telecommunications ecosystem, and the cable operator perspective is being brought to these forums as standards and best practices are developed.

From a network perspective, the operator edge and access networks are beginning to incorporate DAA, which uses Ethernet-based optical transport, 10GbE today, and distributed compute and networking power in unsecured outside plant. Most operators' 10G vision involves DAA. This introduces a different type of security risk level compared to cable-specific analog optics and "dumb" HFC nodes. In addition, proprietary CCAP platforms – not invulnerable but purpose-built and highly customized systems – are giving way to virtual platforms, consisting of off-the-shelf servers and compute for packet processing, again creating a different category of security threat type. CableLabs' Distributed Virtualization Security is focused on threat assessment and recommendations for these NFV-based DAA systems, including understanding how similar problems are being tackled in other sectors moving in the same direction, such as the wireless and telco industries.

Another innovative program underway at CableLabs in support of 10G is the MicroNets project [19]. As the name may imply, MicroNets apply to the home, and aim to harden the home network against the increasing number of devices with a range of integrity. With applications racing ahead and devices multiplying, bad actors are at an advantage as standards for security mature. MicroNets recognizes this range of trustworthiness or hackability, and the basic premise is to create smaller networks, "micronets," within the traditionally single home LAN, with the walls of division associated with the security level of the devices and service. Issues that arise are isolated to a specific MicroNet, ensuring that as evolving IoT security addresses devices over time, adjacent devices and service are not impacted.

There is much more to come on the security front, as both emerging networking for 10G and target applications act to increase the significance of thinking through security up-front in the design development process.

## **7. Conclusion**

Announced merely 21 months ago, 10G has rapidly become the industry's new North Star. As can be observed herein, the extremely efficient name of the initiative describes quite a wide range of technology elements that it is building upon and building anew. There are many reasons to be looking for the reset button for the year 2020. However, with the direct and critical contributions of our very industry, people and the network are adapting. Timelier than ever, progress towards 10G has not lost a step – indeed it is accelerating, as operators are moving from technology assessments and white board sessions to investment decisions, project definition and development (the fun stuff!). A comparison of 10G material from the 2019 SCTE Expo and the 2020 SCTE Expo shows remarkable progress on many fronts. It only magnifies the excitement for what 2021 will bring, masks or no masks!

## Abbreviations

|        |                                                |
|--------|------------------------------------------------|
| ACI    | Adjacent Channel Interference                  |
| AI     | Artificial Intelligence                        |
| CAGR   | Compound Annual Growth Rate                    |
| CCI    | co-channel interference                        |
| CLGD   | cable load generator                           |
| COTS   | Commercial off-the-shelf                       |
| CP     | Cyclic Prefix                                  |
| CPE    | Customer Premises Equipment                    |
| CPD    | Common Path Distortion                         |
| DAA    | Distributed Access Architecture                |
| DAAS   | Distributed Access Architecture Switch         |
| EC     | Echo Cancellation                              |
| FBC    | Full Band Capture                              |
| FDD    | Frequency Domain Duplex                        |
| FDX    | Full Duplex DOCSIS                             |
| FEC    | Forward Error Correction                       |
| Gbps   | Gigabits per second                            |
| H-AGG  | Hub Aggregation                                |
| ICCAP  | Integrated Cable Modem Termination System      |
| Igs    | Interference Groups                            |
| IUC    | Interval Usage Codes                           |
| LDPC   | Low Density Parity Check                       |
| LUL    | Latency Under Load                             |
| M2M    | Machine 2 Machine                              |
| MER    | Modulation Error Ratio                         |
| ML     | Machine Learning                               |
| MMP    | Multiple Modulation Profiles                   |
| OFDM   | Orthogonal Frequency Division Multiplexing     |
| OFDMA  | Orthogonal Frequency Division Multiple Access  |
| PBH    | Peak busy hours                                |
| PIM    | Passive Intermodulation                        |
| PMA    | Profile Management Application                 |
| PNM    | Proactive Network Maintenance                  |
| QAM    | Quadrature Amplitude Modulation                |
| QB     | Queue Building                                 |
| RBAs   | Resource Block Assignments                     |
| RBW    | Resolution Bandwidth                           |
| RP     | Roll-off Period                                |
| SC-QAM | Single Carrier Quadrature Amplitude Modulation |
| SG     | Service Group                                  |
| SID    | Spectral Impairment Detection                  |
| SIK    | Self-install Kit                               |
| SLA    | Service Level Agreement                        |
| TaFDM  | Time and Frequency Division Multiple           |

|     |                            |
|-----|----------------------------|
| TDR | Time Domain Reflectometer  |
| TGs | Transmission Groups        |
| UGS | Unsolicited Grant Services |
| VOP | Velocity of Propagation    |
| VSA | Vector Signal Analyzer     |

## Bibliography & References

- [1] Barker, Bruce E, and Claude Bou Abboud, Erik Neeld, “Access Capacity Planning: Staying Well Ahead of Customer Demand Helped Ensure Stability During COVID-19,” SCTE Cable-Tec Expo, Oct 13-16, 2020.
- [2] Bienkowski, Tom, “Pandemic Life Online: What is Everybody Doing,” [www.netscout.com](http://www.netscout.com).
- [3] Cloonan, Dr Tom, John Ulm, Ayham Al-Banna, Frank O’Keefe, and Ruth Cloonan, “Capacity Planning, Traffic Engineering, And HFC Plant Evolution For The Next 25 Years, SCTE Cable-Tec Expo, Sept 30-Oct 3, 2019, New Orleans, LA.
- [4] Cloonan, Dr. Tom, Ruth Cloonan, and John Ulm, “Updated FDX Traffic Engineering Analysis,” CommScope presentation to Comcast, 7/13/20-8/10/20.
- [5] Dolan, Andy, Maintaining Confidentiality on the 10G Network  
<https://www.cablelabs.com/maintaining-confidentiality-10g-network>
- [6] Howald, Dr. Robert L, Comcast, “Aboard the Technology Wave: Surf Report,” SCTE Cable-Tec Expo, Sept 26-29, 2016 Philadelphia, PA.
- [7] Howald, Dr. Robert L, Comcast, “The Fiber Frontier,” INTX Spring Technical Forum, May 16-18, 2016, Boston, MA.
- [8] Howald, Dr. Robert L, , “Network Preparation: Maximizing Capacity ROI,” SCTE Cable-Tec Expo, Oct 21-24, 2013, Atlanta, GA.
- [9] Howald, Dr. Robert L, and Dr. Sebnem Ozer, Robert Thompson, Saif Rahman, Dr. Richard Prodan, Jorge Salinger, “What is 10G – The Technology Foundation,” SCTE Cable-Tec Expo, Sept 30-Oct 3, 2019, New Orleans, LA.
- [10] Mutalik, Venk, D. Rice, K. Subramanya, J. Wang What Gets Measured Gets Done / What Gets Analyzed Gets Transformed, , <https://www.nctatechnicalpapers.com/Paper/2018/2018-analytics-for-a-wider-deeper-network-view>, SCTE Cable-Tec Expo, Oct 22-25, 2018.
- [11] Nandiraju, Nagesh, “Distributed Access Architecture – Goals and Methods of Virtualizing Cable Access,” SCTE Cable-Tec Expo, Sept 26-29, 2016 Philadelphia, PA.
- [12] Prodan, Dr. Richard, “Optimizing the 10G Transition to Full Duplex DOCSIS 4.0,” SCTE Cable-Tec Expo, Oct 13-16, 2020.

- [13] Rice, Daniel, M. Harb, J. Ferreria, B. Santangelo, R. Spanbauer “A Machine Learning Pipeline for D3.1 Profile Management,” SCTE Cable-Tec Expo, Sept 30-Oct 3, 2019, New Orleans, LA.
- [14] Rice, Daniel, M. Harb, J. Ferreria, B. Santangelo, R. Spanbauer, “Full Scale Deployment of PMA: Lessons Learned from Deploying the Profile Management Application System at Scale and Plans for Expanding the System Beyond the OFDM Scope”, SCTE Cable-Tec Expo, Oct 13-16, 2020.
- [15] Thompson, Robert, C. Moore, J. Moran, R. Howald, Optimizing Upstream Throughput Using Equalization Coefficient Analysis, <https://www.nctatechnicalpapers.com/Paper/2009/2009-optimizing-upstream-throughput-using-equalization-coefficient-analysis>, 2009.
- [16] Wall, Dr. Bill, “Practical Considerations For Full Duplex Deployments In N+x Environments”, SCTE Cable-Tec Expo, Sept 30-Oct 3, 2019, New Orleans, LA
- [17] CableLabs DOCSIS 4.0 PHY Specification: CM-SP-PHYv4.0-I01-190815.docx
- [18] <https://www.cablelabs.com/10g/security>
- [19] <https://www.cablelabs.com/technologies/micronets>

# Smart Data Powers Service Layer Management For Network Operations 2.0

A Technical Paper prepared for SCTE•ISBE by

**Dr Vikram Saksena**

Cable CTO

NetScout Systems

310 Littleton Road, Westford, MA 01886

978-614-4383

vikram.saksena@netscout.com

**Ryan Eccles**

Principal Sales Engineer

NetScout Systems

# Table of Contents

| <b>Title</b>                                                    | <b>Page Number</b> |
|-----------------------------------------------------------------|--------------------|
| 1. Introduction.....                                            | 3                  |
| 2. Extracting Service Intelligence with Smart Data.....         | 3                  |
| 3. Addressing Operational Challenges with Smart Analytics ..... | 4                  |
| 3.1. Proactive Service Assurance .....                          | 4                  |
| 3.2. Enabling Zero-Touch Automation.....                        | 5                  |
| 3.3. Improving Customer Experience .....                        | 5                  |
| 3.4. Early Warning Detection for Advanced Threats .....         | 5                  |
| 3.5. Enabling Just-in-time Resource Management .....            | 5                  |
| 4. Smart Data Use Cases in Operator WiFi Networks.....          | 5                  |
| 4.1. Instrumentation.....                                       | 6                  |
| 4.2. Use Case 1: Slow Customer Login .....                      | 6                  |
| 4.3. Use Case 2: Intermittent Customer Attach Problems.....     | 7                  |
| 4.4. Use Case 3: Decreasing WiFi Network Usage .....            | 8                  |
| 5. Conclusion .....                                             | 9                  |
| Abbreviations.....                                              | 9                  |

## List of Figures

| <b>Title</b>                                                       | <b>Page Number</b> |
|--------------------------------------------------------------------|--------------------|
| Figure 1 - High Definition Smart Data.....                         | 3                  |
| Figure 2 - Smart Analytics.....                                    | 4                  |
| Figure 3 - WiFi Instrumentation Points .....                       | 6                  |
| Figure 4 - Login Latency over Time.....                            | 7                  |
| Figure 5 - Vanity SSID Success/Failed Transactions over Time ..... | 7                  |
| Figure 6 - Vanity SSID Portal Service Monitor.....                 | 8                  |
| Figure 7 - Packets Showing Incomplete DHCP Exchange .....          | 8                  |
| Figure 8 - Custom Smart Data Dashboard .....                       | 9                  |



# 1. Introduction

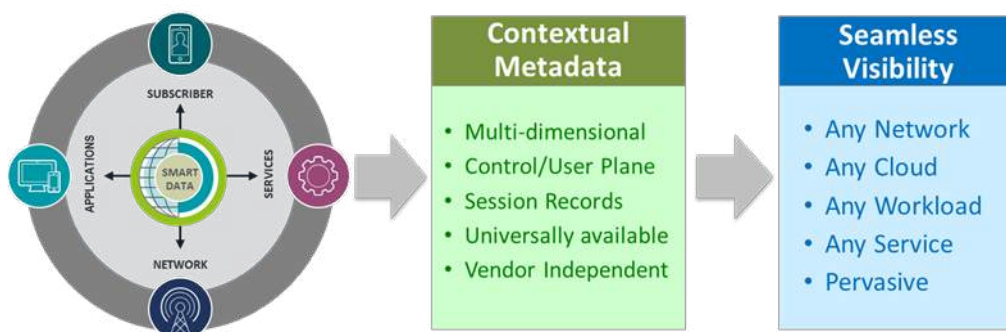
Cable operators have embarked on a network transformation journey to evolve their networks from a centrally controlled, purpose-built platform to a distributed, programmable platform. A programmable network fabric drives service agility by allowing operators to create, change, and personalize services on demand, thereby shrinking service delivery intervals significantly. Together with business process automation and new service innovation, operators are well positioned to drive revenue growth and profitability. In this new era of service agility, operators will have to raise the bar on differentiating themselves on superior service quality to retain and grow their customer base amidst a highly dynamic and competitive landscape. In this paper we describe a powerful approach to service layer management utilizing packet data.

As operators evolve to a distributed, service-oriented architecture, the delineation between network and service management becomes more pronounced. Machine data (syslogs, telemetry, flows, session records, etc.) collected from the network elements is better suited for network management as they directly point to problems in the underlying infrastructure. Being vendor specific and fragmented, it is challenging to correlate machine data from various elements in the network to get an accurate view into service layer problems. A better and more direct way to identify service layer problems is to look at packet data, which is vendor independent, universally available, and carries all the contextual and relevant information pertaining to subscribers and the services they consume. Every transaction between subscribers or between a subscriber and an application leaves a footprint on the operator network which can be analyzed for service performance issues as well as security related threats.

Smart Data, which is actionable metadata derived from packets, delivers superior service management capabilities due to its high-definition nature. Smart Data coupled with Smart Analytics enables proactive service assurance, network automation, just-in-time resource management, SLA management and many other use cases at the speed and agility required for Network Operations 2.0. We also describe an actual operator use case for proactive service assurance of Wi-Fi networks.

## 2. Extracting Service Intelligence with Smart Data

Smart Data is high-definition, actionable intelligence derived from the ultimate source of truth, IP packets. It is contextual metadata that is structured, timely, and relevant. Smart Data gives operators complete visibility with holistic, end-to-end coverage of control and user plane traffic over multiple dimensions of their network, services, and subscribers. It addresses user experience covering all devices, network services, and applications consumed.

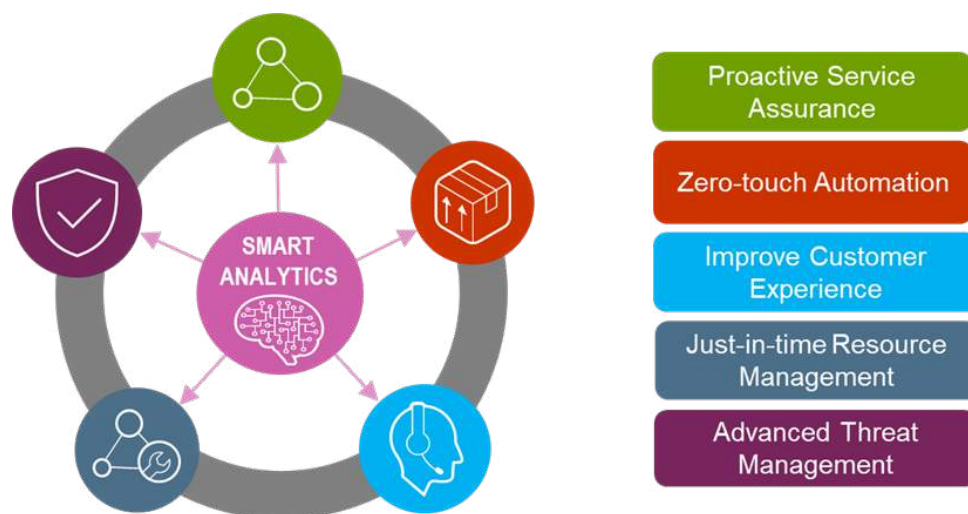


**Figure 1 - High Definition Smart Data**

Smart Data lowers the cost/complexity barrier of packet-based instrumentation and makes it affordable for pervasive edge to core deployment. Being a software solution it has an optimized, elastic footprint that scales up and down with traffic levels. It offers vendor independence by providing unified metadata with common views and workflows for any service (voice, data, video, IoT), over any network (physical, virtual, fixed or mobile), over any cloud (private or public), and for any workload (virtual-machine or container-based).

### 3. Addressing Operational Challenges with Smart Analytics

Smart Data is the fuel for Smart Analytics that addresses many operational challenges in current and future networks: providing proactive service assurance, enabling zero-touch automation, reducing churn and improving net promoter score, enabling just-in-time resource management, and advanced threat management for security operations. Having a unified data model that can be leveraged across functional and organization boundaries helps accelerate digital transformation by breaking down the information silos.



**Figure 2 - Smart Analytics**

#### 3.1. Proactive Service Assurance

As services become mission-critical, especially for commercial customers, proactive service assurance becomes very important. Smart Data enables service assurance for proactive alerting, rapid problem isolation, and situational analysis. It helps in promoting trouble-free network operation by identifying potential problems before they become service-affecting.

For example, as cable operators deploy mobile services over a combination of Wi-Fi and cellular technologies (4G/5G), the complexity for end to end service assurance increases. Handovers across Wi-Fi and cellular technology domains must be carefully monitored to avoid dropped calls and degraded sessions. From an end user perspective, the service experience must be consistent and disruption-free across both network domains.

### **3.2. Enable Zero-Touch Automation**

One of the benefits of moving to a software-centric network is that corrective actions in the network can be automated with little or no human intervention. In coordination with policy servers and orchestration platforms, service performance triggers based on Smart Data can be used to effect network changes in near real-time. Additional resources can be temporarily increased for virtual machines and containers to deal with traffic spikes in the network by the appropriate use of service performance triggers fed to the policy and orchestration platforms. Such “closed-loop” automation solutions can drastically reduce or eliminate service degradations that could negatively impact user experience.

### **3.3. Improve Customer Experience**

Creating a composite customer experience indicator from Smart Data helps operators better understand their customer experience for the services they consume. Each user is measured for key performance indicators related to accessibility, retainability, and quality of experience for deriving the indicator. With such a customer experience indicator it is possible to visualize the experience for groups of users in different dimensions such as device type, location, service and more. When this indicator falls below the desired levels, network operations can drill down to the details of what caused the degradation. Understanding and rapidly responding to network causes for poor service experience is an important factor in improving customer experience and reducing churn.

### **3.4. Enable Just-in-time Resource Management**

In a hardware-centric network, capacity management required careful planning because hardware had to be ordered and deployed ahead of realized demand. This process created inefficiencies and stranded capital in the network when demand would materialize differently than what was planned for. In a software-centric network, traffic data can be used to scale resources up or down in relatively shorter intervals than what was done before. Analytics based on Smart Data can be used to deploy just-in-time resources as dictated by key service performance indicators. Such a dynamic resource management process allows for more efficient use of capital as capacity deployments can be made to track more closely with the realized traffic demands.

### **3.5. Early Warning Detection of Advanced Threats**

Security threats continue to rise. As networks evolve to a software centric architecture, new attack surfaces are exposed to potential hackers. Hackers have become more sophisticated and continue to develop new attack vectors for exploiting vulnerabilities in open source software and virtualized, cloud-native infrastructures. Smart Data provides deep packet intelligence to identify the potential onslaught of advanced threats which would otherwise escape detection. Sophisticated security analytics algorithms based on Smart Data provide early warning detection of anomalous behavior and alert operators for taking proactive actions to avoid service disruptions.

## **4. Smart Data Use Cases in Operator WiFi Networks**

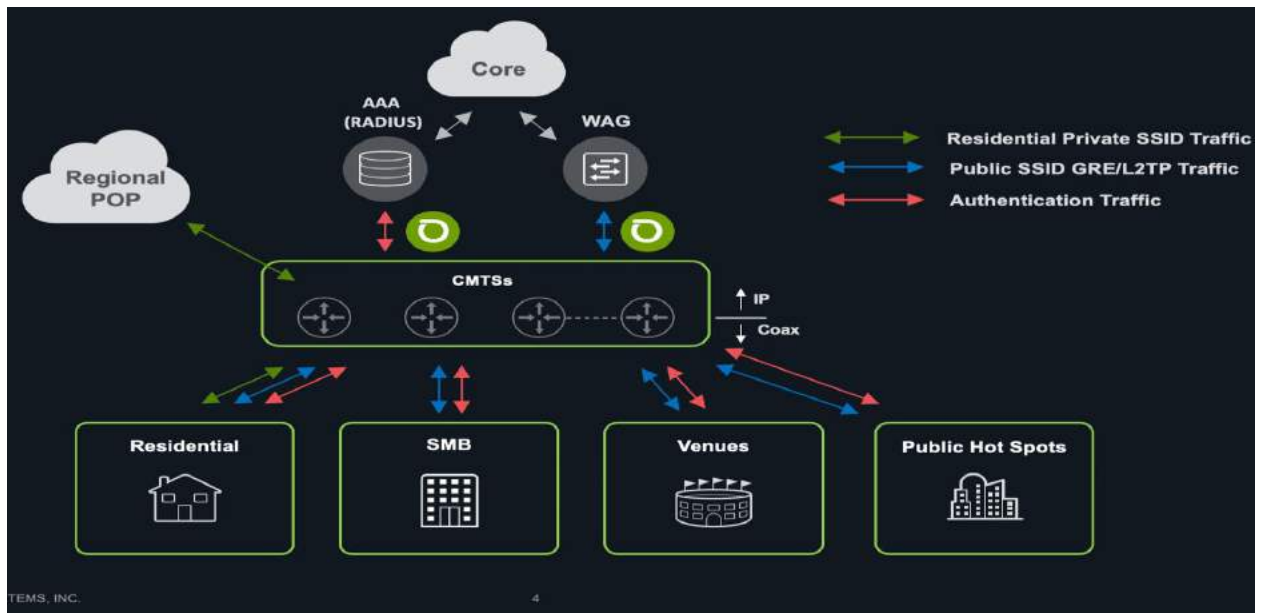
Wi-Fi networks are now held to a much higher standard than ever before. No longer is best effort wireless service considered acceptable. Additionally, the proliferation of MVNO agreements is further requiring the Wi-Fi service to be especially robust in order to meet subscriber experience requirements. Some of the common challenges that Cable MSOs face in offering carrier-grade Wi-Fi service include:

- Identifying customer login experience problems
- Problem isolation with limited subscriber feedback
- Identifying Wireless Access Gateway problems
- Last mile customer experience telemetry
- Controller stability
- Determination of Wi-Fi performance for operator provided devices versus other devices

This paper will focus on the first three use cases and will demonstrate how the use of Smart Data can help operators become more proactive, isolate the source of problems and reduce service disruptions.

#### 4.1. Instrumentation

Proper instrumentation and Smart Data generation in modern operator Wi-Fi networks requires tapping and monitoring of the links that surround critical application flows that make up the wireless service offering. Figure 3 below shows the common visibility points (indicated by the NetScout logo) in an operator Wi-Fi environment for maximum visibility.




**Figure 3 - Wi-Fi Instrumentation Points**

#### 4.2. Use Case 1: Slow Customer Login

In many public carrier Wi-Fi environments there is an inherent lack of customer service calls from the subscriber base even during a systemic outage. Subscribers traditionally have grown to view public Wi-Fi as best effort. This is changing rapidly. The ability to use packet-based Smart Data is more important than ever to solving problems quickly and identifying the true subscriber experience. Packet-based metrics are the ultimate source of truth and play a critical role in reducing churn and allowing the operator to become truly proactive.

Figure 4 below illustrates how Smart Data can identify and report on how long it takes for a mobile device to attach and obtain an IP address from the operator Wi-Fi network. The login latency distribution metrics can be dimensioned based on market, region, physical location, AP vendor or even client MAC address.




**Figure 4 - Login Latency over Time**

The figure above illustrates that the majority of devices will get an IP address within the 1 second to 1.5 second range. The ability to trend such data over time and alert on deviations from baseline values is critical to assuring the Wi-Fi service and improving subscriber experience, further illustrating the value of what a Smart Data model can provide in such scenarios.

### 4.3. Use Case 2: Intermittent Customer Attach Problems

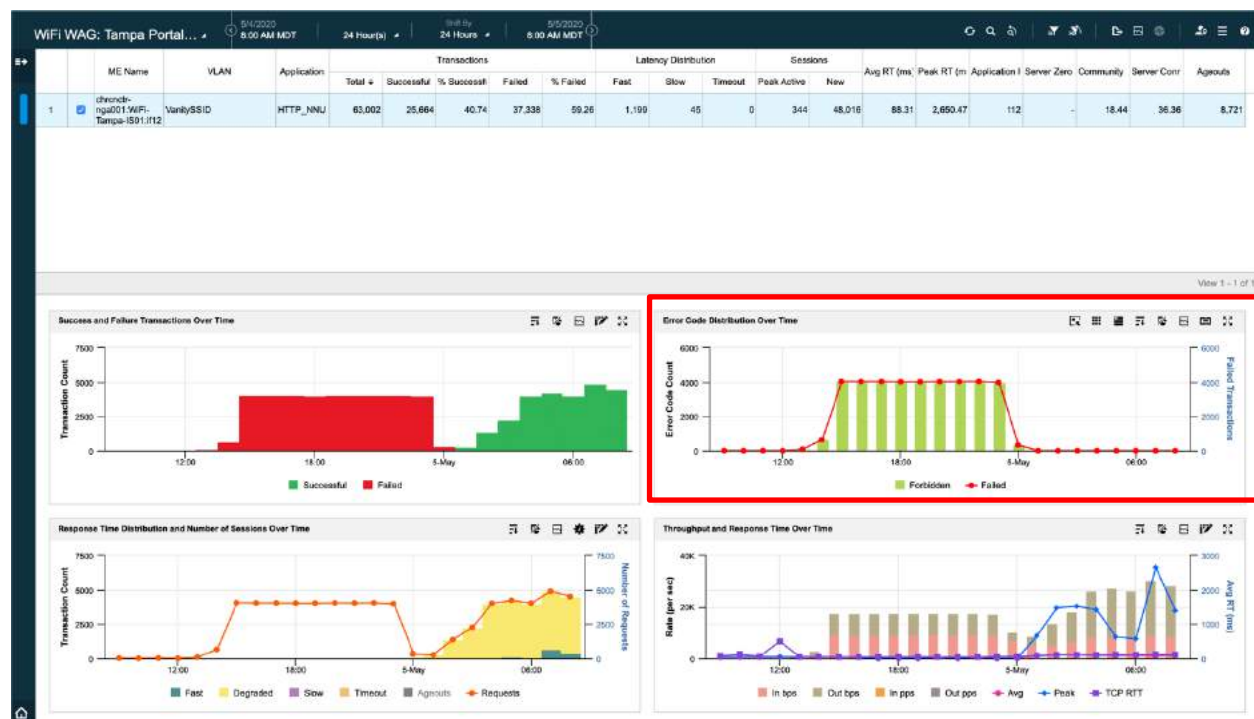
An MSO provider of Wi-Fi service had a low number of problem calls where business subscribers would state that their customers could not connect to their network while waiting for service. The number of calls was not highly significant but was enough to cause the operator to look into whether or not this was a systemic problem.

Using multi-dimensional Smart Data, the operator was able to look back in time at only business vanity SSIDs and trend the portal performance. Figure 5 below shows a 31-day service dashboard view of captive portal transactions for business SSIDs.



**Figure 5 - Vanity SSID Success/Failed Transactions over Time**

The operator found that there were periods of low success rates with one being the most severe. Figure 6 below shows a service monitor drill down into Vanity SSID portal transactions for that particular day.



**Figure 6 - Vanity SSID Portal Service Monitor**

The service monitor shows that during the middle of the day there was a spike in HTTP 403: Forbidden errors. Further drilldown into the packet-based Smart Data showed that the devices that should have been granted access to the Wi-Fi network were not. This data allowed the operator to examine recent provisioning changes that were done incorrectly and put proper measures in place.

#### 4.4. Use Case 3: Decreasing WiFi Network Usage

An MSO offering Wi-Fi service began to detect a downturn in usage over a period of weeks when the opposite was expected. Multiple systems including AAA were checked but nothing explaining the downturn was detected in the data that those systems produced. As a result the MSO turned to packet-based Smart Data to provide a view into individual device behavior.

| No. | Time        | Source               | Destination     | Protocol | Length | Info                                                                     |
|-----|-------------|----------------------|-----------------|----------|--------|--------------------------------------------------------------------------|
| 1   | 0.00000000  |                      |                 | XID      | 72     | Basic Format; Type 1 LLC (Class I LLC); Window Size 0                    |
| 2   | 0.076308680 | 0.0.0.0              | 255.255.255.255 | DHCP     | 388    | DHCP Discover - Transaction ID 0x29bea3d4                                |
| 3   | 0.076397520 | 100.64.250.1         | 100.64.250.250  | DHCP     | 366    | DHCP Offer - Transaction ID 0x29bea3d4                                   |
| 4   | 0.081758660 |                      |                 | ICMPv6   | 132    | Neighbor Solicitation for                                                |
| 5   | 0.081763640 |                      |                 | ICMPv6   | 108    | Router Solicitation                                                      |
| 6   | 1.054609880 |                      |                 | ICMPv6   | 156    | Multicast Listener Report Message v2[Packet size limited during capture] |
| 7   | 1.171126700 | 0.0.0.0              | 255.255.255.255 | DHCP     | 388    | DHCP Request - Transaction ID 0x29bea3d4                                 |
| 8   | 1.171220190 | 100.64.250.1         | 100.64.250.250  | DHCP     | 378    | DHCP ACK - Transaction ID 0x29bea3d4                                     |
| 9   | 1.575724530 | 0.0.0.0              | 255.255.255.255 | DHCP     | 388    | DHCP Discover - Transaction ID 0x29bea3d5                                |
| 10  | 1.575811560 | 100.64.250.1         | 100.64.250.250  | DHCP     | 366    | DHCP Offer - Transaction ID 0x29bea3d5                                   |
| 11  | 2.064557190 | fe80::14a9:b772:673_ | ff02::16        | ICMPv6   | 156    | Multicast Listener Report Message v2[Packet size limited during capture] |

**Figure 7 - Packets Showing Incomplete DHCP Exchange**

The packets in Figure 7 above revealed that there were devices that were not fully completing the IP acquisition process with the Wireless Access Gateway.



The raw packets were sent to the Wireless Access Gateway vendor and it was identified that there was a bug in their software that would prevent devices from completing the IP acquisition process in specific circumstances. A custom dashboard shown in Figure 8 was then produced to identify the same issue going forward.

| Client Equipment MAC Address | SSID | DHCP Sessions | Zero Downlink Throughput | Zero Uplink Throughput |
|------------------------------|------|---------------|--------------------------|------------------------|
| 76:00:00:00:00:00            |      | 0 sessions    | 0.00 B                   | 7,232.00 B             |
| 96:00:00:00:00:00            |      | 1 sessions    | 0.00 B                   | 0.00 B                 |
| 00:00:00:00:00:00            | WiFi | 1 sessions    | 0.00 B                   | 0.00 B                 |
| f4:00:00:00:00:00            | WiFi | 1 sessions    | 0.00 B                   | 0.00 B                 |
| 08:00:00:00:00:00            |      | 0 sessions    | 0.00 B                   | 13,528.00 B            |
| 44:00:00:00:00:00            |      | 5 sessions    | 0.00 B                   | 0.00 B                 |
| dc:00:00:00:00:00            | WiFi | 1 sessions    | 0.00 B                   | 0.00 B                 |
| 50:00:00:00:00:00            |      | 2 sessions    | 0.00 B                   | 0.00 B                 |
| c0:00:00:00:00:00            |      | 1 sessions    | 0.00 B                   | 0.00 B                 |
| 00:00:00:00:00:00            |      | 1 sessions    | 0.00 B                   | 0.00 B                 |

**Figure 8 - Custom Smart Data Dashboard**

## 5. Conclusion

Smart Data technology utilizes packet data and generates elastic metadata through pervasive, edge-to-core instrumentation in a distributed operator network. This metadata can be consumed by a variety of upstream applications focused on service operations, network automation, customer experience, SLA management, and security. By leveraging Smart Data technology, the operator teams can proactively monitor, assure and secure all IP services (voice, data, and video) which may run across different domains (physical, virtual, or cloud) and access networks (fixed or wireless) using the same views and workflows. By relying on a common, high definition data platform that provides “visibility without borders”, operators can go through their network transformation journey with confidence to innovate while delivering an unparalleled customer experience.

## Abbreviations

|      |                                 |
|------|---------------------------------|
| MVNO | Mobile Virtual Network Operator |
| MAC  | Media Access Control            |
| SSID | Service Set Identifier          |
| HTTP | Hyper Text Transfer Protocol    |
| IoT  | Internet of Things              |
| SLA  | Service Level Agreement         |
| MSO  | Multi System Operator           |

# **Constructing a Convergence Lab**

## **Lessons Learned From Building a Converged Network at CableLabs**

A Technical Paper prepared for SCTE•ISBE by

**Matthew Schmitt**  
Principal Architect  
CableLabs  
Pleasanton, CA  
303-661-9100  
[m.schmitt@cablelabs.com](mailto:m.schmitt@cablelabs.com)



# Table of Contents

| Title                                                        | Page Number |
|--------------------------------------------------------------|-------------|
| 1. Introduction.....                                         | 3           |
| 2. DAA and P2P Coherent Optics Background.....               | 3           |
| 2.1. Pre-DAA Hybrid Fiber-Coax (HFC) Networks.....           | 3           |
| 2.2. A Fundamental Shift in Cable Access Networks.....       | 4           |
| 2.3. The Need for Aggregation.....                           | 5           |
| 2.4. The Case for Coherent.....                              | 5           |
| 2.5. The Pathway to a Converged Network Infrastructure ..... | 6           |
| 3. The Plan.....                                             | 7           |
| 3.1. Why Build a Convergence Lab.....                        | 7           |
| 3.2. Finding Partners.....                                   | 8           |
| 3.3. The Construction Plan.....                              | 10          |
| 3.4. The Response .....                                      | 10          |
| 4. The New Reality .....                                     | 11          |
| 4.1. Restricted Access .....                                 | 11          |
| 4.2. Equipment Arrivals .....                                | 11          |
| 4.3. A New Hope .....                                        | 13          |
| 4.4. Everything takes longer in a pandemic .....             | 14          |
| 5. Today and Tomorrow .....                                  | 17          |
| 5.1. Current Status .....                                    | 17          |
| 5.2. Next Steps.....                                         | 19          |
| 5.3. A note of thanks .....                                  | 20          |
| 6. Conclusion.....                                           | 21          |
| Abbreviations .....                                          | 21          |

## List of Figures

| Title                                                         | Page Number |
|---------------------------------------------------------------|-------------|
| Figure 1 – Simplified Cable HFC Architecture .....            | 3           |
| Figure 2 – Simplified Distributed Access Architecture .....   | 4           |
| Figure 3 – DAA with Multiple Child Nodes .....                | 5           |
| Figure 4 – Converged Network Infrastructure.....              | 7           |
| Figure 5 – Phase 1: P2P Coherent Optics.....                  | 8           |
| Figure 6 – Phase 2: Distributed CCAP Architecture .....       | 9           |
| Figure 7 – Phase 3: Wireless.....                             | 9           |
| Figure 8 – Phase 4: PON .....                                 | 10          |
| Figure 9 – Packages piling up.....                            | 12          |
| Figure 10 – Tower of boxes .....                              | 13          |
| Figure 11 – 220V power .....                                  | 15          |
| Figure 12 – The plug doesn't fit .....                        | 16          |
| Figure 13 – Ciena to Ciena communication.....                 | 18          |
| Figure 14 – EXFO traffic generators with Ciena switches ..... | 18          |
| Figure 15 – ADVA switch with EXFO traffic generator .....     | 19          |

# 1. Introduction

Cable operators are in the midst of deploying or preparing to deploy a variety of distributed access architecture (DAA) approaches. When this is done by aggregating multiple child nodes together onto a single point-to-point (P2P) coherent optics link, it has the effect of pushing very high capacity fiber-based Ethernet links deep into their networks. Beyond improving their DOCSIS®-based residential broadband services over coax, it also opens up numerous new business opportunities—such as mobile fronthaul and backhaul, business ethernet, remote passive optical networks (PON) and more—along with the ability to converge all of these different technologies onto a single network infrastructure.

CableLabs® has been working jointly with our members and technology vendors to develop the tools and technologies needed to make this vision a reality. However, while it's one thing to develop the technology, it's quite another to build devices based on it, and still another to integrate it together into a working whole. While CableLabs is not in a position to manufacture those devices, we are well positioned to bring them all together in a new network infrastructure convergence lab.


This convergence lab will allow us to demonstrate how these various technologies can be brought together to operate over a single physical network, showcasing the work that's already been done, and learning the implications of operating them together. It will also create a platform on which to test out new technologies going forward to see how they integrate together into a converged whole.

This paper goes through some of the background on DAA and P2P coherent optics technology, reviews the process we went through to build the lab, any lessons learned thus far, and the current status of what is an ongoing process of continuing to grow and develop this lab.

## 2. DAA and P2P Coherent Optics Background

### 2.1. Pre-DAA Hybrid Fiber-Coax (HFC) Networks

The typical access network architecture deployed by cable operators prior to the advent of DAA approaches employed a Hybrid Fiber-Coax (HFC) approach to reach their customers with video and data services. A simplified version of this architecture is shown in Figure 1 below.



**Figure 1 – Simplified Cable HFC Architecture**

In this example, downstream radio frequency (RF) signals for data and video are generated in a Hub facility, converted from electrical to optical using an analog laser, transported over fiber to a Fiber Node at which point the signal is converted back to electrical, which is then transmitted over a coaxial cable network to end customers. In the upstream a similar set of operations occurs, originating at the customer premises and terminating at the Hub. The distance from the Hub to the Fiber Node is typically 20-80 km,


although longer distances are possible; the distance from the Fiber Node to the end customer is generally just a few km.

When the capacity requirements of the group of customers sharing a portion of the network (a Service Group) exceeds the capacity of a Fiber Node, operators will typically split a Fiber Node into multiple Nodes. Each of those Fiber Nodes will require a separate link back to the Hub, but will often share a single fiber pair using Dense Wavelength Division Multiplexing (DWDM) to save costs. Each new Service Group also requires new equipment at the Hub site, which can create pressure on space, power, and cooling.

## 2.2. A Fundamental Shift in Cable Access Networks

In response to these pressures, as has been covered in previous papers and presentations, many cable operators have begun a fundamental shift in their networks by moving the RF generation previously performed in the Hub out into the network. This allows them to convert their existing fiber network from analog optics to digital optics and to utilize Ethernet networking across the fiber portion of their cable plant. The effect is to create a deep fiber Ethernet network.

This approach is what we refer to as a Distributed Access Architecture (DAA), because functions that were previously centralized are now distributed. A simplified version is shown in Figure 2 below. On the surface it looks very similar to Figure 1: homes are still connected to coax, which terminates at a device that is connected to fiber, which in turn is connected to the hub. However, under the surface things are dramatically different: the fiber is carrying Ethernet data over digital optics rather than modulated RF signals over analog optics, and the RF is generated at a Remote PHY Device (RPD) or Remote MAC/PHY Device (RMD) directly connected to the coax without the electrical-optical-electrical conversion that occurred previously.



**Figure 2 – Simplified Distributed Access Architecture**

This change has a number of advantages:


- It reduces the equipment requirements for Hubs;
- It allows cable operators to utilize lower cost off-the-shelf digital optics for the fiber link instead of specialized analog optics; and
- By avoiding electrical-optical-electrical conversions, the signal to noise ratio (SNR) of the electrical signal on the coaxial cable is significantly improved.

That SNR improvement allows devices based on the DOCSIS 3.1 specifications to operate at higher modulation orders, which increases total network capacity.

### 2.3. The Need for Aggregation

Each RPD or RMD typically has a 10 gigabit-per-second (Gbps) port, which provides sufficient capacity for generating a full spectrum of digital video and DOCSIS 3.1 signals. Therefore, if Fiber Nodes were simply being replaced with RPDs or RMDs on a 1:1 basis, all that would be required to connect them to the Hub would be a single 10 Gbps optical link.

However, a key driver for the move to DAA approaches is to enable more Service Groups with fewer customers on each one. Figure 3 shows a simplified version of a scenario where the single service group from Figures 1 and 2 is split into multiple smaller service groups, each serviced by a separate RPD or RMD (each of which we refer to as a child node), and aggregated together at an aggregation node (AN) onto the existing fiber. The AN sits where the Fiber Node was previously.



**Figure 3 – DAA with Multiple Child Nodes**

One straightforward means of building this approach would be to passively aggregate 10 Gbps optical links from each child node using DWDM, with the AN containing a passive mux/demux. This technology is readily available and well understood, and so represents a logical initial approach.

However, it has its limits, since at a typical spacing of 100 GHz for each signal cable operators will be limited to about 48 of these 10 Gbps links over a single fiber pair. And, the cost goes up linearly with every 10 Gbps link that is added.

### 2.4. The Case for Coherent

Architects at CableLabs postulated that utilizing a P2P coherent optics link operating at 100 Gbps or 200 Gbps per wavelength might prove advantageous over an approach using multiple 10 Gbps optical links with DWDM as described above. One end of the coherent optics link would be terminated at the Hub facility, and the other end at a device located in the AN called a Coherent Termination Device (CTD); the CTD would also terminate multiple 10 Gbps links from the Child Nodes, which could utilize low cost grey optics operating at 10 km or less. The aggregation could be performed in multiple different ways: at layer 1 with a muxponder; at layer 2 with a switch; or at layer 3 with a router.

Conceptually this looks exactly the same as what was shown in Figure 3 above, with the following exceptions: the AN has a CTD in it which forwards traffic at layer 1, 2, or 3 rather than containing a DWDM mux; the links connecting the CTD to the Child Nodes use low cost grey optics (instead of higher cost colored optics); and instead of carrying those 10 Gbps signals all the way back to the Hub (which requires longer reach, higher cost optics), they are aggregated onto one or more P2P coherent optics links for transport back to the Hub.

CableLabs staff conducted an analysis to compare the costs associated with each approach for supporting multiple child nodes in the same footprint as an existing fiber node. We found that the key determination for which approach would cost less depended on the number of child nodes. If the number of child nodes being aggregated together is relatively small, the DWDM approach was more cost effective, because the cost of the P2P coherent optics transceiver and the CTD is shared across only a few child nodes. However, as the number of child nodes increases that dynamic flips, because the cost of the DWDM approach goes up linearly with each child node due to the need to add a pair of higher-cost tunable 10 Gbps transceivers, whereas with the P2P coherent optics approach the cost for each additional child node is much smaller because only very low cost grey optics transceivers are added.

Further, while there are already cases where the P2P coherent optics approach was lower cost, we also saw potential opportunities to bring that cost down further, which would lower that crossover point.

Using P2P coherent optics provides additional benefits as well: it can coexist with existing signals using DWDM, which in turn also opens the door to add additional P2P coherent optics links in the future. The net result is a 10-20x increase in available capacity for the fiber that cable operators have already deployed into the range of terabits per second, extending the lifetime of their existing network investment dramatically.

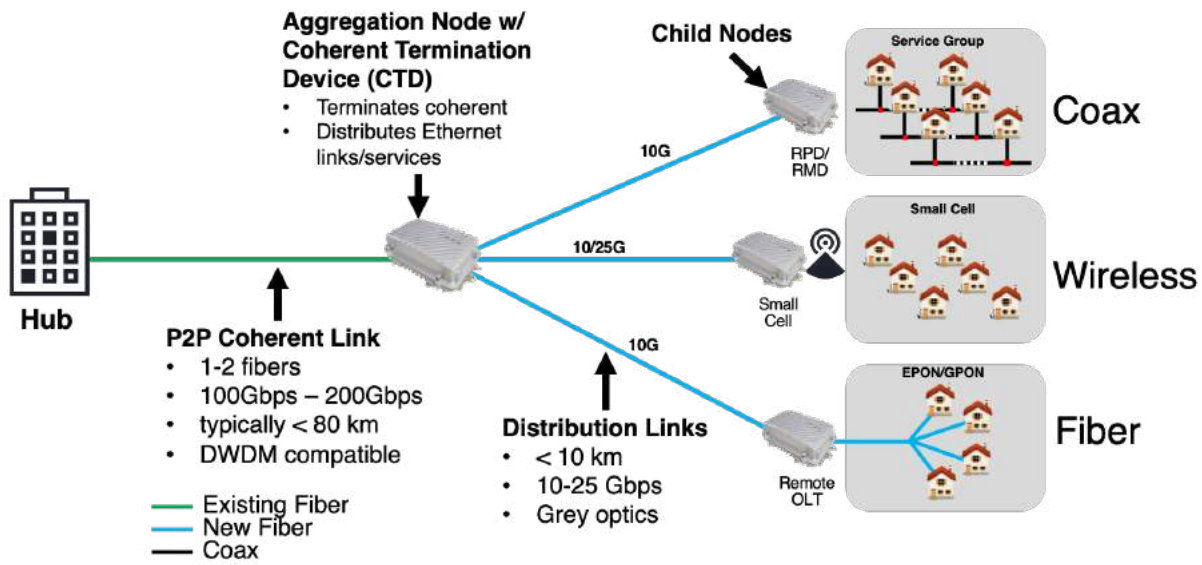
## **2.5. The Pathway to a Converged Network Infrastructure**

In addition to the benefits identified above, moving to an architecture utilizing P2P coherent optics opens the door to new business opportunities by enabling a converged network infrastructure.

A DOCSIS network is, in effect, a very deep Ethernet network optimized for running over an HFC infrastructure. However, the only points at which you could access that network are at cable modems attached to the coax portion of the plant: accessing the fiber portion of the network cannot be done via DOCSIS, and instead was only possible by using separate fibers or separate wavelengths over the existing fiber.

However, let's assume you've deployed a P2P coherent optics architecture using a network switch in the CTD. That CTD has become an Ethernet connection point that can tap directly into very high capacity fiber links. Any service that can operate with an Ethernet connection can connect at that point, including DOCSIS services, PON services, mobile xhaul services, business services, etc. All of them can be carried over the same Ethernet based fiber network, aggregated together via the CTD.

The result is a converged network infrastructure, as shown in Figure 4 below.



**Figure 4 – Converged Network Infrastructure**

With this type of architecture approach, the opportunities are limited only by our imaginations.

In order to help enable this future, CableLabs worked with our members and manufacturers to develop a series of specifications that define requirements for coherent optics transceivers operating at 100G and 200G. The intent was to provide manufacturers with guidance regarding the requirements for operating on a cable operator network, as well as to promote interoperability, which in turn helps promote scale and competition, thereby reducing cost. The specification for 100G per wavelength operation was initially released in 2018, and the specification for 200G per wavelength operation was initially released in 2019.

### 3. The Plan

#### 3.1. Why Build a Convergence Lab

Another advantage of P2P coherent optics is that there is existing equipment that can be leveraged now, with more on the way. Interoperable 100G coherent optics transceivers that are compliant to the CableLabs specification requirements already exist, as do a variety of switches, routers, and muxponders that they can operate in, some of which are temperature hardened. Prototypes of interoperable 200G transceivers are expected this year as well.

However, we realized that a key piece – a fully weather hardened Aggregation Node with a CTD that could operate in it – was missing.

The plant architectures utilized by many cable operators around the world favor a clamshell type device that is fully weather hardened at the AN location. Unfortunately, that device doesn't yet exist and needs to be developed. In order for manufacturers to decide to invest the money to do that, they need to understand what requirements their customers have, and confirm that there's sufficient interest from those customers.

One way in which CableLabs is working to address these needs is by working with our members to define common requirements for a CTD.



Another way in which we realized we might be able to help was by demonstrating the capabilities of this new architecture, showing not only that it can support existing use cases exceptionally well, but even more importantly that it can support a broad range of other applications and services that could provide new revenue opportunities.

Thus, the idea for building a network infrastructure convergence lab was born. It would be a platform on which we could demonstrate how a variety of different services could coexist: DOCSIS broadband over coax; mobile fronthaul and backhaul; FTTH via PON; and business ethernet services. A showcase for products built to comply with a wide range of CableLabs specifications. And a test bed for validating how new products and technologies could integrate into that same infrastructure.

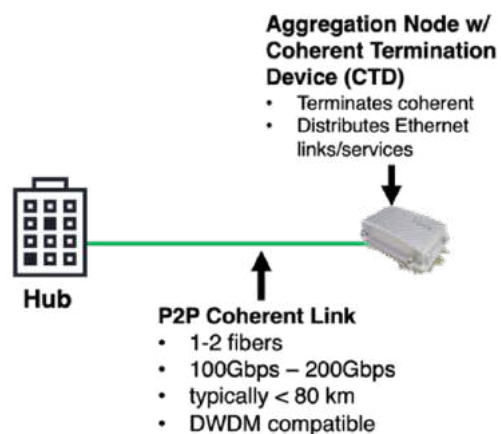
### 3.2. Finding Partners

In order to make that a reality, a key requirement would be obtaining the equipment—ideally from a variety of manufacturers—in order to make it work. So began a quest to find partners in this endeavor.

Utilizing relationships we had developed in putting together the CableLabs P2P Coherent Optics specifications, as well as relationships that were developed as a part of building the P2P Coherent Optics interoperability demo shown at SCTE Expo 2019, and also relationships from the work done with CableLabs on Distributed Access Architectures (DAA), we reached out to a number of manufacturers with a fairly simple pitch to build the network shown in Figure 4 above in a lab at CableLabs. The benefit for manufacturers: feedback regarding how their devices operate in the network, visibility to our members when we report on the activity, and visibility to anyone that visits CableLabs and tours our lab spaces.

As a part of these discussions, we proposed to develop the lab in four phases.

The first phase would be focused on the main optical link from the “Hub” to the CTD, since all other applications and services would run over that infrastructure.

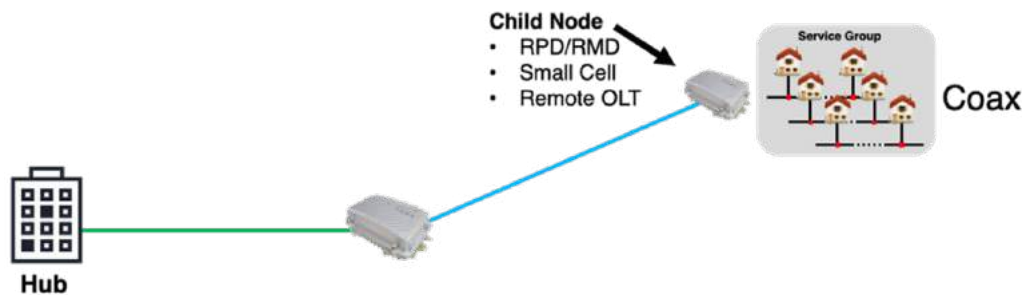


**Figure 5 – Phase 1: P2P Coherent Optics**

Ideally it would include equipment from multiple manufacturers so that we could demonstrate interoperability, and multiple transceivers operating simultaneously. That would demonstrate the ability of the solution to scale up as needed with multiple wavelengths operating at the same time. Playing the role of the CTD in the AN would be temperature-hardened equipment that would be suitable for deployment in a street cabinet (since clamshell devices were not available).

As a part of this phase we would also connect traffic generators directly to the CTD in the AN as well as in the Hub, which serves the dual purposes of validating the link and showing how a business ethernet service might work.

The second phase would be focused on building a Distributed Converged Cable Access Platform (CCAP) Architecture (DCA) network that would operate over the top of the P2P coherent optics link serving as the backbone of the architecture, since this would likely be the first, primary application for cable operators.



**Figure 6 – Phase 2: Distributed CCAP Architecture**

The intent was to start with the Remote PHY Devices (RPDs) that cable operators are starting to deploy, working with a CCAP Core in the Hub location.

The third phase would then look at layering in wireless solutions, which present an immense new business opportunity for cable operators, particularly with the rollout of 5G services that—in fronthaul applications in particular—have high capacity and low latency requirements, while also requiring higher densities than are typically the case today due to the use of higher frequency spectrum.




**Figure 7 – Phase 3: Wireless**

Within this, there would be opportunities to demonstrate the use of both direct fiber connections as well as DOCSIS links for backhaul and fronthaul transport of mobile traffic.

The fourth phase would incorporate a remote PON solution, with the remote OLT being either a standalone device connected via a short fiber run to the CTD (as shown below) or a device that connects directly to the CTD (such as an OLT in an SFP+ form factor).





**Figure 8 – Phase 4: PON**

Note that any of the follow on phases after Phase 1 could be done in any particular order depending on equipment availability, and additional phases to add equipment to address other use cases would be possible.

Therefore, our primary focus was on finding partners to supply equipment to support Phase 1, with a lesser focus on Phase 2, leaving the remaining phases as future work.

### **3.3. The Construction Plan**

Another question we addressed during these discussions was how to go about actually constructing the lab setup.

As mentioned, at SCTE Expo in 2019 the CableLabs booth featured a P2P coherent optics interoperability demo which included contributions from a number of different manufacturers. To prepare for that demo, we had held a sort of mini-interop, where each of the companies providing equipment sent an engineer out to CableLabs, and we spent several days piecing the demo together and ensuring we could make it work. We then took the demo apart, shipped it to New Orleans, and rebuilt the same setup we had used in the lab with substantial help from some of those same contributors.

For this effort, we proposed to do essentially the same thing: have each of the participating companies send out an engineer to work together over the space of several days to install the equipment and get it all working. In some ways it would be even easier than what we did the previous year, because we wouldn't have to tear it down and ship it out to be rebuilt; we could build it right where it stay and be used. If any issues came up, we could deal with them on the spot.

It seemed a straightforward, easy approach for getting the lab up and running quickly.

### **3.4. The Response**

After making our pitch, a number of manufacturers agreed to provide equipment either on loan, as a donation, or at a substantial discount. Specifically, the following list of companies (in alphabetical order) offered the following pieces of equipment:

- II-VI (formerly Finisar): a C Form-factor Pluggable 2 – Analog Coherent Optic (CFP2-ACO) transceiver module
- ADVA Optical: a network switch that supports coherent optics, targeted for the “Aggregation Node” location, along with coherent optics transceivers
- Ciena: two network switches that support coherent optics, one each for the “Hub” and “Aggregation Node” locations, along with coherent optics transceivers

- Edge-Core: a whitebox network switch that supports coherent optics, targeted for the “Hub” location, including support for CFP2-ACO modules
- EXFO: a pair of traffic generators capable at running multiple 100G streams of data simultaneously
- Lumentum: a CFP2-ACO coherent optics transceiver module
- Vecima: a Remote-PHY Device (RPD) and associated management station

With verbal agreements in place, we began working on all the necessary paperwork to allow these transactions to take place, making arrangements for equipment to be shipped out, etc.

All was going relatively smoothly, if perhaps a bit slower than we’d all have preferred: getting executive approvals, legal approvals, and documents signed always seems to take longer than you’d expect. Still, given that the actual construction of the lab should only take a matter of days with engineering support on site, some amount of delay wouldn’t prevent us from completing the lab in a reasonable amount of time. Enough so that I felt comfortable submitting a request to SCTE to create this paper.

Then the Covid-19 pandemic hit with full force, blowing up many of our plans.

## **4. The New Reality**

### **4.1. Restricted Access**

Like many companies, as the rapid spread of the virus became clear, CableLabs instituted a work from home policy. Specifically, unless being in the building was absolutely essential to the core operation of the company, we were all to work from home. Further, travel between our California and Colorado offices was halted, and visitor access to the building was cut off.

This obviously presented several problems for the lab construction plan. First, I’m located in California, but the lab was to be built at our main office in Colorado; therefore, I would be unable to work on building the lab myself. Second—and most significantly—we would not be able to have any outside engineers visit the facility, meaning we would not be able to rely on their expertise to get their equipment operational. And finally, since this work didn’t necessarily count as essential to the core operation of the company, even staff located in Colorado were unable to enter the building to work on the lab.

However, at that point in time, there was still hope that these restrictions would only be temporary: after a month or two of a tight lockdown things would get better and we’d be able to travel and access the labs again. Besides which, it was taking longer than expected to get the equipment agreements signed, equipment shipped, etc. So, we had time to spare and hoped that by waiting things out we would still be able to build the lab as planned (albeit perhaps a bit later than originally envisioned).

Therefore, we continued to work on arranging equipment shipments to our facility in Colorado so that it would be ready to go as soon as virus restrictions were lifted.

### **4.2. Equipment Arrivals**

Thanks in no small part to the contributions and efforts of the various companies that had partnered with us on this effort, who had continued to work through the pandemic and their own restrictions, the needed equipment did indeed start to arrive. However, since there was no one there to receive, unpack, and work with the equipment it began to pile up, as shown in Figure 9 below.



**Figure 9 – Packages piling up**

In fact, as shown above, the boxes were starting to clog up one of our hallways. As soon as one of our local team members was able to re-enter the office, we found a more efficient stacking arrangement, as shown in Figure 10.



**Figure 10 – Tower of boxes**

So now we had much of the equipment we needed, but still no way to put things together until our level of access to the building changed.

### **4.3. A New Hope**

While the situation with the pandemic did not improve nearly as much or as quickly as we might have hoped, it did improve enough that CableLabs was able to relax some of our building access restrictions. Visitors to Colorado were still out of the question, but local staff would be permitted limited access to the building for short periods of time when arranged in advance.

Among those that would be given some access to the building were the members of the CableLabs Optical Center for Excellence, the team that we have working on next generation optical technologies: because they can't bring their experiments home, they needed access to the optical lab in order to continue their work. If they could spare some time to lend a hand, given their expertise, they would be perfect for setting up the lab.

Fortunately for me, they are some of the most helpful, generous people you'd ever want to meet, and they agreed to help me out and work on setting up the equipment there on site without complaint. The caveat was that they also had to make progress on their own work, and so would have only a limited amount of time each week to help out. But I was grateful for any help I could get.

A wave of outreach to our partners who were supplying the equipment to request remote support to our on-site team followed; they were all supportive, and so a series of introductions were made to connect those that would be supporting the work remotely with the team on the ground.

Our prospects for building the lab were looking up.

#### **4.4. Everything takes longer in a pandemic**

The local team was able to access the lab starting in mid-July. They began by unpacking equipment from one of our partners and trying to get it setup and running with remote engineering support.

What I didn't yet recognize—but would soon be brought face-to-face with—was that everything takes longer in a pandemic. In particular, when you can only access the lab two days a week (in order to limit exposure), and are relying on remote support, problems that might've been resolved in hours or a day if people were on-site and dedicated to the task can take a week or more.

And it seems like the simplest problems are the ones that catch you off-guard and cause the most problems.

For example, we discovered when we unpacked a network switch from one of our partners that it required a 220V power supply, something we didn't expect—since all the equipment we'd worked with up to this point worked on 110V—but admittedly should have asked about in advance. After several calls and emails, we discovered that we did in fact have 220V power supplies in the building at CableLabs, and members of our Kyrio lab team were able to set one up for us.



**Figure 11 – 220V power**

We initially thought we had overcome that issue with relative ease and minimal impact, but quickly realized that we weren't going to get off that easy: the power supply we had at CableLabs used an L6-20 receptacle (rated for up to 20 amps), whereas the plug for the box we were trying to install and setup used an L6-30 plug (rated for up to 30 amps).



**Figure 12 – The plug doesn't fit**

We inquired with our partner regarding an alternate power cord, but were told that what we had was the only one they provided. We asked if the box would ever draw more than 20 amps, and were told that it probably wouldn't, although they didn't want to commit to that at that time. As a result, no one was comfortable simply using an adapter or replacing a plug. Instead, we started to look into purchasing a new 30 amp power dongle for our rail system, which we found would take at least 6 weeks.

This left us at a bit of a stand-still on this particular device until we eventually made two critical discoveries: while the plug on one end of the power cord that came with this network switch was specifically a 30 amp plug, the connector on the other end of the power cord was in fact only rated for 20 amps; and buried deep in the technical specs of the network switch we were using, we discovered that the switch itself had a max draw of just 9.3 amps. With that information in hand we became comfortable with the idea of using an adapter to connect the power cord to our 20 amp power supply, and quickly ordered one. And while that took just a few days to arrive, it also had to wait for the next window when someone would be in the lab, further delaying the day by which we could power up that device.

As I said, under these unusual circumstances, the simplest things become multi-week delays.

The power plug issue was just one example; each device setup ran into their own simple yet time consuming issues.

For one of them, the issue was licensing. For this particular manufacturer, they assumed that devices would always be connected to a live internet/network connection. Therefore, they had been designed to pull their license from over the internet, without which it wasn't operational.

However, in our case, we were setting these up in a lab without an external network connection; and while that might eventually happen, it certainly was not the case for our initial setup. There was an option to create a download for us to use, although we'd have to do it through their customer portal in which we

weren't set up. It didn't get setup correctly the first time around, given it was a non-standard way of doing things. In the end we got it all working, but only after multiple emails, phone calls, and aborted attempts.

For another network switch we were attempting to setup, the issue wasn't power or licensing, it was transceiver compatibility.

This partner had provided us with a network switch under the assumption—which we shared—that one of the coherent optics modules we already had on hand would work. And initially it did, as we were able to power up the switch and have it recognize the transceiver module. All seemed fine. However, we then discovered that we couldn't set the frequency on which we wanted it to operate. It would only operate on a single fixed frequency, even though we knew the module was capable of tuning to different frequencies. Eventually the problem was traced to a compatibility issue: the switch would only support controlling the coherent transceiver module if it were also from the same manufacturer.

Our partner who had supplied the switch graciously offered to provide us with a pair of loaner transceivers, and moved to make that happen as quickly as possible. However, quickly in this case means getting a loaner agreement drawn up, having it reviewed by one legal department, having those changes reviewed by another, and finally getting it signed and processed. And then of course, waiting for shipping.

All of our issues were ones that could be overcome given time, but all of which also took longer simply because of the unique situation we found ourselves in with the pandemic.


## **5. Today and Tomorrow**

### **5.1. Current Status**

As of the writing of this text, here's where we stand today.

We now have both of the two Ciena switches up and running in our lab, with a transceiver in each unit talking to a transceiver in the other, as shown in Figure 13 below.






**Figure 13 – Ciena to Ciena communication**

The EXFO FTB-4 Pro traffic generators were recently unpacked and are now up and running; we're currently working on getting them setup to pass traffic across each of our switches.



**Figure 14 – EXFO traffic generators with Ciena switches**

We've also gotten the ADVA switch up and running with in a loopback configuration.



**Figure 15 – ADVA switch with EXFO traffic generator**

Other equipment is yet to be unboxed but will be soon.

## **5.2. Next Steps**

In the short term, our plan is to get traffic running across the Ciena and ADVA switches, unpack and get the Edge-Core Cassini running with the CFP2-ACO modules that were provided to us by II-VI and Lumentum, and connect them all together using DWDM.

That will form the backbone of our lab to demonstrate network infrastructure convergence.

From there we'll be working on getting a Remote PHY implementation running over top of our optical network using the equipment provided by Vecima—in combination with a CCAP-Core already on-site—to demonstrate that service. That'll be followed by other phases as we're able, establishing that you can run multiple services over the same access network infrastructure simultaneously.

Beyond that, we plan to use this lab as both a showcase (it'll eventually be moved to a more visible location) and a working lab that we can use to test and demonstrate new components of a converged network. For example, we'll be interfacing it with equipment from other teams within CableLabs, such as those working on network virtualization, to create a truly converged network of the future. It's going to take a while to get there—much longer than I ever would have anticipated—but many of the building blocks are already in place, and more will be soon.

### 5.3. A note of thanks

While I may be the one writing this paper, given the restrictions on travel and building/lab access, I haven't been the one doing the work on the ground as I had hoped and intended. Instead, that's been done by a team of folks from multiple companies to whom I am greatly indebted:

- CableLabs
  - Haipeng Zhang
  - Jon Schnoor
  - Junwen Zhang
  - Mu Xu
  - Scott Bybee
  - Steve Jia
- Kyrio
  - Efren Torres
  - Mark Ambrozic
- ADVA
  - Clark Scott
  - Jack Yocum
  - Morgan Shaner
- Ciena
  - Craig McCoy
  - Darren McKinney
  - Fernando Villarruel
- Edge-Core
  - Jeff Catlin
- II-VI
  - John DeAndrea
  - Shan Esser
- IP Infusion
  - Srikanth Krishnamohan
- Lumentum
  - Ed Kirchoff
  - Ernest Muhigana
  - Jan Willem Poelman
  - Scott Swail
- Vecima
  - Colin Howlett
  - James Trueman

My sincere thanks to all of you, without whom none of what we have done (and will be doing!) would have been possible. And to anyone whom I may have inadvertently left off of this list, my apologies and deepest thanks.

This has been very much an industry wide effort, and it is my heartfelt hope that it will continue to be one going forward.

## 6. Conclusion

When I originally proposed this paper, the expectation was that I would be writing about what we had learned by successfully building a converged network infrastructure: that I would be describing why we had set out to build it, what had happened when we did, and hopefully about how we had demonstrated that these various services could operate effectively over a common access network infrastructure, along with perhaps a few issues the MSOs should watch out for in their own networks.

And while it did accomplish the first of those objectives, from there it instead evolved into a story regarding the very unusual and unique circumstances in which we all find ourselves, and the challenges that imposed in this particular situation. And a story about how we can push through those challenges when we all work together to overcome them.

Which to me is one of the key lessons to be learned here: that while everything takes longer in a pandemic, when we work as a team, we can push through those challenges and accomplish a lot.

This network infrastructure convergence lab is a work in progress, and in fact I hope and expect it will always remain one, with new equipment and technologies being incorporated for years to come. In fact, additional contributions are more than welcome, as we'd love to enhance what we can do in the lab and bring in as many different pieces of equipment for as many different services from as many different manufacturers as possible.

The lab may not be as far along as we'd like, but thanks to the contributions of multiple individuals we will get there, which I hope to be able to report on in the future.

## Abbreviations

|          |                                                   |
|----------|---------------------------------------------------|
| AN       | Aggregation node                                  |
| CFP2-ACO | C Form-factor pluggable 2 – analog coherent optic |
| CCAP     | Converged cable access platform                   |
| CTD      | Coherent termination device                       |
| DAA      | Distributed access architecture                   |
| DWDM     | Dense wave division multiplexing                  |
| G / Gbps | Gigabits per second                               |
| HFC      | hybrid fiber-coax                                 |
| P2P      | Point-to-point                                    |
| PON      | Passive optical network                           |
| RF       | Radio frequency                                   |
| RPD      | Remote PHY device                                 |
| RMD      | Remote MAC/PHY device                             |

# **Using Machine Learning To Automate Node Split Designs And HFC Augmentation Options**

A Technical Paper prepared for SCTE•ISBE by

**Keith R. Hayes**  
Chief Executive Officer  
IMMCO, Inc.  
12395 Morris Rd, Alpharetta, GA 30005  
770-378-3595  
Keith.hayes@immcoinc.com

# Table of Contents

| Title                                                                      | Page Number |
|----------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                       | 3           |
| 2. Artificial Intelligence (AI) and Machine Learning (ML).....             | 3           |
| 3. Network Capacity Augmentation Methods.....                              | 4           |
| 4. Implementing Machine Learning for Node Split Design.....                | 5           |
| 4.1. Programming Environment .....                                         | 5           |
| 4.1.1. Programming Language .....                                          | 5           |
| 4.1.2. IDE – Integrated Development Engine .....                           | 5           |
| 4.1.3. Learning Process .....                                              | 6           |
| 4.1.4. Decision Tree / Classification Engine .....                         | 6           |
| 4.2. Network Data Extract .....                                            | 6           |
| 4.3. Business Rules / Design Rules .....                                   | 7           |
| 5. Let's go split a node....with Machine Learning! .....                   | 7           |
| 5.1. Machine Learning Node Split Design Process.....                       | 7           |
| 5.2. Implementing and Scaling Node Split Machine Learning Environment..... | 11          |
| 6. Creating a Holistic Network Capacity Augmentation Environment .....     | 12          |
| 7. Conclusion.....                                                         | 13          |
| Abbreviations .....                                                        | 13          |
| Bibliography & References.....                                             | 13          |

## List of Figures

| Title                                                                                 | Page Number |
|---------------------------------------------------------------------------------------|-------------|
| Figure 1 – Relationship of Artificial Intelligence and Machine Learning .....         | 4           |
| Figure 2 – Subset of Network Data Extract.....                                        | 6           |
| Figure 3 – ML Network Element Relational Schematic .....                              | 8           |
| Figure 4 – Network Elements near the original Node .....                              | 8           |
| Figure 5 – New Node Location .....                                                    | 9           |
| Figure 6 – HHP after minimal construction Node Split .....                            | 9           |
| Figure 7 – ML – Designed Node Split with balanced HHP .....                           | 10          |
| Figure 8 – ML-designed Node Split HHP data.....                                       | 11          |
| Figure 9 – Possible inputs to holistic HFC Analysis, Design, and Planning Engine..... | 12          |

# 1. Introduction

- **>300,000** Nodes
- **~700,000** Power Supplies
- **~2,000,000** HFC plant miles
- **>70,000,000** US MSO Internet Consumers
- **>40% CAGR** Broadband Consumption
- **Streaming video** here and growing, **online game-streaming** and **AR/VR** coming
- Oh, by the way, **COVID-19** induced **telework** and **virtual classroom** data activity

There are more than 300,000 HFC nodes in the US currently, and several million more worldwide. The additional network traffic triggered by Covid-19 in the spring of 2020 increased the level of HFC capacity augments by as much as 300% compared to 2019 volume. Network augmentation techniques such as node splits, adding HSI EIA's (6 MHz channels), service group de-combines, bandwidth expansion, Node+0/RPD's and mid-/high split reverse path expansion require engineering and operations resources in both ISP and OSP.

This paper will examine techniques in which much of this activity can be automated and iteratively optimized through Machine Learning (ML). With inputs provided from network mapping systems, capacity monitoring platforms, spectrum management applications, and business rules such as preferred augmentation hierarchy, expected duration before next augment, balancing of house-counts, municipal permitting difficulty, and cost efficiency, the ML environment would rapidly analyze entire geographic segments of the network, and provide augmentation planning data including network design changes and BOM's for areas requiring immediate physical layer upgrades.


As the ML platform iteratively processes network geographies, it will learn from how past predictions tracked to current status and continuously adjust to optimize capacity augmentation methods and designs. As the ML environment will be analyzing the entire network geography, data to drive Capex planning for future years will be derived enabling the operator to more efficiently allocate capital.

The ML environment would also support "what/if" network topology planning for approaches, such as bandwidth expansion to 1.8 GHz vs Node+0 at current bandwidth, and provide data on cost, duration before next augment needed, and percentage of network elements requiring replacement or repositioning.

## 2. Artificial Intelligence (AI) and Machine Learning (ML)

There is a lot of interest in using advanced computing technology to automate and accelerate processes that when done by humans are slow and tedious – facial recognition for example.

AI and ML are often used interchangeably but as depicted in the figure below ML is a subset of AI – so how are they related?



**Figure 1 – Relationship of Artificial Intelligence and Machine Learning**

**Artificial Intelligence (AI)** can best be understood as developing computer systems to perform tasks that humans can typically do better. A mechanical example would be a robotic welder in an auto manufacturing plant. A data analytics example would be Recommendations from your online retailer or streaming service that takes your past history and analyzes it against current inventory to “Recommend” items you may want to purchase or content you may want to watch.

**Machine Learning (ML)** is a subset of AI, and in addition to enabling computer systems to perform tasks with humans can typically do better, deploys “learning algorithms” that allows the system to measure performance and automatically improve outputs from those learnings. Back to one of the previous examples, if the robotic welder was coupled with an x-ray or gamma-ray measurement system that provided data on weld quality which then would change the parameters of the welding programming to improve performance without human intervention would be a good example of ML.

### **3. Network Capacity Augmentation Methods**

One of the unique benefits of HFC architecture is the many options potentially available to add capacity, with more being developed as technology advances. The list below is not exhaustive but notes most of the typical options to add capacity without changing the bandwidth or US/DS ratio of the network, organized by rough order of difficulty to implement:

- Activate unused / repurpose EIA’s (channels)
- Service Group de-combining/recombining
- Increase modulation density
- Extend fiber to heavy-use households
- Node Splits
  - Add transceivers to existing node (segmentation)
  - Add new node – minimum construction, no HHP balancing



- Add new node – balance HHP
- Add new node – balance for peak data utilization
- Add Remote Phy-Mac/Phy device

As evaluations are performed to determine the best approach on a node-by-node basis many other variables come into play:

- Municipal Permitting
- Maintenance Window availability
- Aerial/Underground

These many methods and variable are further complicated by the situation that most operators have separate departments managing ISP (Inside Plant) and OSP (Outside Plant) requiring careful coordination for those methods requiring both groups.

## **4. Implementing Machine Learning for Node Split Design**

### **4.1. Programming Environment**

When developing a Machine Learning application there are several fundamental environment decisions that must be made:

- Which programming language?
- Which IDE? (Integrated Development Environment)
- Which learning process – Supervised or Unsupervised?
- Which Decision Tree / Classification Engine?

#### **4.1.1. Programming Language**

Most common programming languages are capable of being easily used for ML environments, including:

- Python
- Java
- C++
- C#
- R
- JavaScript
- Scala

#### **4.1.2. IDE – Integrated Development Engine**

In order to provide an optimized development and debugging container, an IDE provides numerous tools and applications optimized for ML code development and testing. Many IDE's work with multiple programming languages but not with all, so analysis will need to be done to ensure the chosen programming language and IDE are compatible. Common IDE's include:

- PyCharm
- RStudio
- R-Brain
- Jupyter

- Spyder
- Geany

### 4.1.3. Learning Process

Machine Learning is either Supervised or Unsupervised. Supervised Learning is when the data being processed has been labeled and the ML system is taught by example, such as labeling pictures of items like a hammer and a banana so give the ML environment information to learn from. If the ML algorithm incorrectly identifies a shape, it would be “taught” by human input to improve its accuracy. Unsupervised learning is most useful in clustering and identifying similar objects and detecting significant anomalies – such as fraud detection identifying an abnormal financial transaction base on your spending history.

### 4.1.4 Decision Tree / Classification Engine

The last area that is needed to support the ML environment is a decision tree or classification engine. Common classifiers include:

- GBM Gradient Boosting Machine
- Random Forests
- Logistic Regression

The Python/Jupyter combination with Supervised Learning via random forest classifier was used for the development environment in which to build and test node split design via Machine Learning.

## 4.2. Network Data Extract

To provide data for the ML engine to examine it was necessary to extract every network element, both passives and actives, along with the HHP downstream/upstream, from the network map platform. As each element was geocoded distance calculations could be made without having to extract strand/cable/conduit data simplifying the extraction process. A table showing some of the parameters and results from a design iteration is shown below in Figure 2.

| Equipment      | Element ID | DS HHP | Leg Balancing DS | DS HHP Ratio % | Leg Balancing US | US HHP | US HHP Ratio % | Cascade Limit | Power Supply Proximity | Voltage | Current | Proper Signal Strength | Minimal Construction | Room at Pole or Ped? |
|----------------|------------|--------|------------------|----------------|------------------|--------|----------------|---------------|------------------------|---------|---------|------------------------|----------------------|----------------------|
| BTD-75SH AGC   | 5          | 523    | NO               | 1              | NO               | 0      | 0              | YES           | YES                    | 50.56   | 1.5     | YES                    | YES                  | YES                  |
| BLE-75SH       | 6          | 140    | NO               | 26.8%          | NO               | 383    | 73.2%          | YES           | NO                     | 39.53   | 10.59   | YES                    | YES                  | YES                  |
| 9-TFC-4        |            | 133    | NO               | 25.4%          | NO               | 390    | 74.6%          | YES           | NO                     | 35.58   | 9.5     | YES                    | YES                  | NO                   |
| MB-75SH FD AGC | 10         | 35     | NO               | 6.7%           | NO               | 488    | 93.3%          | YES           | NO                     | 31.92   | 2.91    | NO                     | YES                  | YES                  |
| MB-75SH FD     | 11         | 19     | NO               | 3.6%           | NO               | 504    | 96.4%          | NO            | NO                     | 30.37   | 1.43    | YES                    | YES                  | YES                  |
| BLE-75SH AGC   | 7          | 18     | NO               | 3.4%           | NO               | 505    | 96.6%          | NO            | NO                     | 32.77   | 2.45    | NO                     | YES                  | YES                  |
| MB-75SH FD     | 8          | 8      | NO               | 1.5%           | NO               | 515    | 98.5%          | NO            | NO                     | 31.35   | 1.43    | YES                    | YES                  | YES                  |
| MB-75SH FD     | 9          | 79     | NO               | 15.1%          | NO               | 444    | 84.9%          | YES           | NO                     | 30.24   | 4.39    | NO                     | YES                  | YES                  |
| BLE-75SH       | 68         | 47     | NO               | 9.0%           | NO               | 476    | 89.5%          | NO            | NO                     | 28.89   | 0.99    | NO                     | YES                  | YES                  |
| BLE-75SH       | 70         | 15     | NO               | 2.9%           | NO               | 508    | 95.5%          | NO            | NO                     | 26.83   | 0.99    | YES                    | YES                  | YES                  |
| BLE-75SH       | 69         | 16     | NO               | 3.1%           | NO               | 507    | 95.3%          | NO            | NO                     | 27.32   | 0.99    | YES                    | YES                  | YES                  |
| BLE-75SH       | 12         | 121    | NO               | 23.1%          | NO               | 401    | 75.4%          | NO            | YES                    | 46.52   | 4.89    | YES                    | YES                  | YES                  |
| BLE-75SH AGC   | 13         | 109    | NO               | 20.8%          | NO               | 414    | 79.2%          | NO            | NO                     | 41.15   | 4.15    | YES                    | YES                  | YES                  |
| MB-75SH FD     | 16         | 53     | NO               | 10.1%          | NO               | 470    | 89.9%          | NO            | NO                     | 37.95   | 1.25    | YES                    | YES                  | YES                  |
| BLE-75SH       | 15         | 10     | NO               | 1.9%           | NO               | 513    | 98.1%          | NO            | NO                     | 38.37   | 0.86    | YES                    | YES                  | YES                  |
| MB-75SH FD     | 14         | 34     | NO               | 6.5%           | NO               | 489    | 93.5%          | NO            | NO                     | 39.72   | 1.2     | YES                    | YES                  | YES                  |
| BTD-75SH AGC   | 21         | 203    | YES              | 38.8%          | YES              | 320    | 61.2%          | YES           | YES                    | 51.75   | 4.59    | YES                    | YES                  | YES                  |

**Figure 2 – Subset of Network Data Extract**

Each element from the node to a line extender to a terminating tap was assigned a unique ID number and as the several performance parameters such as signal level, voltage, current and the ratio of HHP both

downstream and upstream from the network element were calculated. The node in the example above had 667 discrete network elements that were loaded into the ML environment.

### **4.3. Business Rules / Design Rules**

After the network data was loaded Business or Design rules had to be created to guide the ML engine based on customer requirements. For example, an operator might want to design for the absolute least amount of activity needed to add capacity – adding another transceiver to segment a node without concern about balancing the homes-passed by the old node vs the new node. Another operator (or heck, maybe the same operator in another part of the network) might want to design for balancing the homes on the old and new node, and provisioning for future nodes. Some of the more than a dozen Rules included in this model were:

- Can the node be segmented?
- Is the proposed location geographically centered or not?
- Are homes passed balanced?
- Is cascade limit exceeded?
- Is there space on the pole or in the pedestal for the node?
- Does RF signal meet specifications?
- Does node meet minimal construction parameters?
- Is any coax reversing needed?
- Is Fiber Splice location reachable?

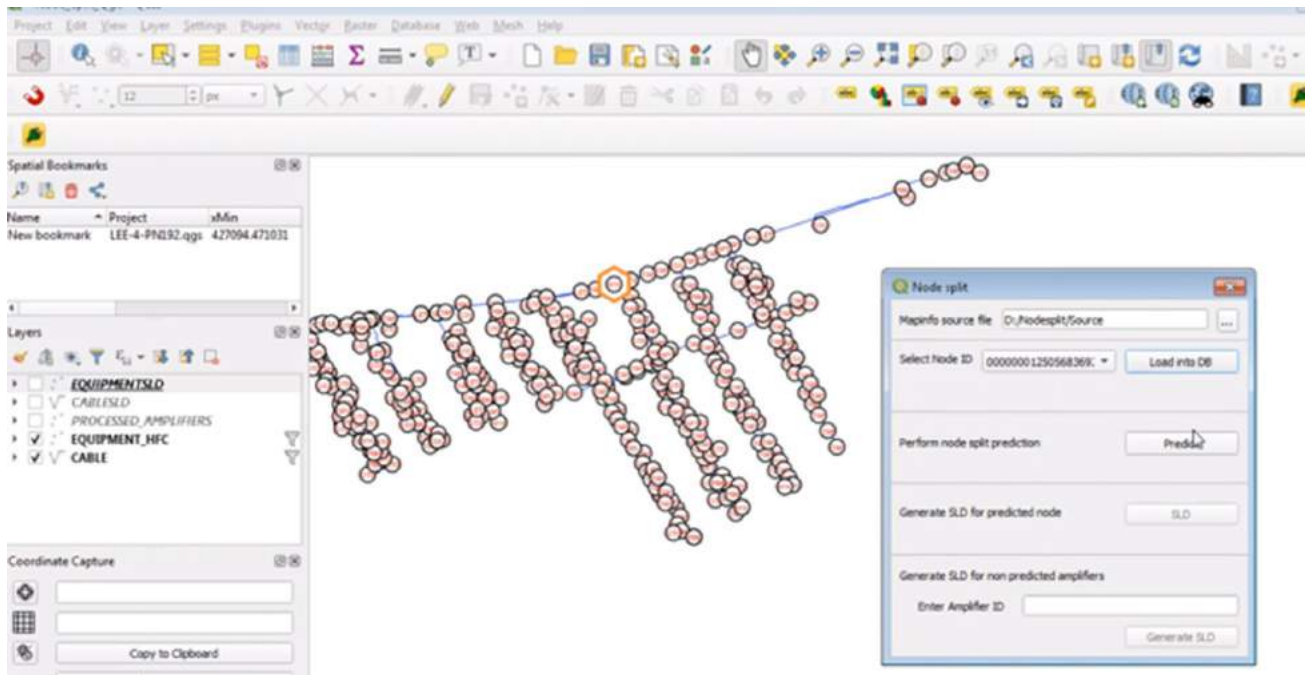
One of the benefits of creating an ML environment is it is relatively easy to add additional parameters and rules as technology changes, for example adding a rule to validate is power ampacity for small cell deployment was available, or tying the network data to the capacity monitoring system which could enable a split location to be balanced on peak data consumption versus geographic center or balanced Homes Passed.

## **5. Let's go split a node....with Machine Learning!**

### **5.1. Machine Learning Node Split Design Process**

In a standard Node Split design process, a designer looks at the network topology, and applying the Business Rules using his or her experience picks a likely location for the new node. Afterwards, calculations have to be completed using a design tool such as Lode Data to verify the location will “work” – all signal levels are acceptable, there is appropriate voltage and no excessive current draw, and if those tests are passed the design needs to be completed and a BOM created. If the new node location does not “work”, the process has to be repeated, each iteration taking 30 minutes or so for an experienced designer. The ML environment was developed to automate this process.

To test the new process several relatively simple node splits were loaded and processed by the ML tool, which generated network element relational schematics as shown in Figure 3 below:



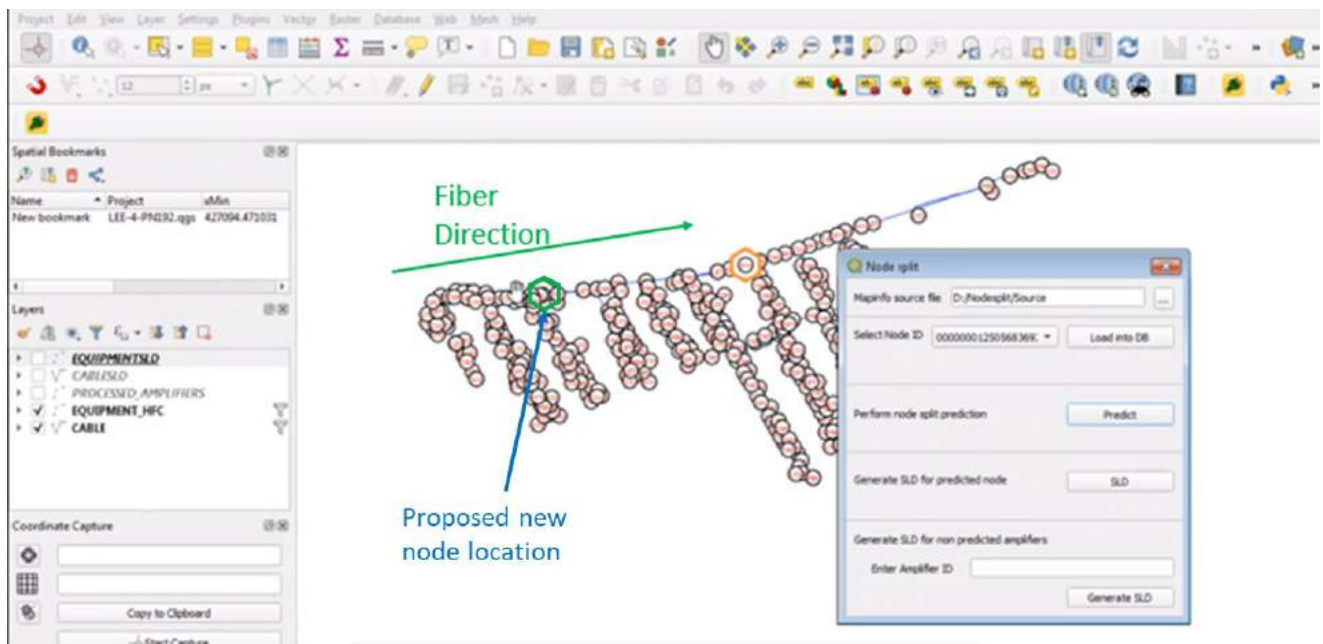
**Figure 3 – ML Network Element Relational Schematic**

The network element schematic is not needed directly by the ML tool but enables a designer to visually review the topology without using a network map system. Figure 4 is a close up showing the original node location and elements near it.



**Figure 4 – Network Elements near the original Node**

This particular Node was non-segmentable, and the Business Rules required minimal construction with no requirement to balance homes passed. The first iteration proposed the new node be placed on the fiber path to the existing node using spare fibers to connect it. As the new node location was an existing amplifier, there was no coax-resplicing required further complying with the minimal construction instruction – Figure 5 below:



**Figure 5 – New Node Location**

When the households passed were calculated for the Original Node and New Node, there was definitely a mis-balance as shown in the table in Figure 6 below, but the operator would be able to provide additional capacity with little cost, time, or likely municipal permitting.


|                           | HHP | Ratio |
|---------------------------|-----|-------|
| Original Node             | 457 | 100%  |
| Original Node after split | 329 | 72%   |
| New Node after split      | 128 | 28%   |

**Figure 6 – HHP after minimal construction Node Split**

To “train” the ML environment, a dozen nodes of varying complexity, size and geographic orientation were used. Each node was “designed” manually with optimum location and business rule adherence, and several deliberately incorrect designs were created for each node configuration. The ML engine “learned” by comparing the correctly designed node split to ones that were incorrect.

After several iterations of supervised feedback, the ML node split algorithm was tested in 200 locations and determined the optimum location in 196 – an accuracy rate of 98%! Even more noteworthy was each location was processed in seconds versus around a half-hour for a human designer.


Below in Figure 7 is an example of a geographically complex node where the ML business rules required HHP balancing, future node segmentation capability, and construction costs and difficulty were not considered.



**Figure 7 – ML – Designed Node Split with balanced HHP**

This node served 523 HHP before the split, and after the split the ML design determined four legs each of which could be segmented into a node in the future with no construction needed, however more than 2,000 feet of fiber and extensive coax-reversing would be required. The HHP data is in the table in Figure 8 below:





|               |     |
|---------------|-----|
| Original Node | 523 |
|---------------|-----|

|              | Original Node<br>Leg A | Original<br>Node Leg B | New Node<br>Leg A | New Node<br>Leg B | Coax<br>Reversing? | Fiber Extension Footage |
|--------------|------------------------|------------------------|-------------------|-------------------|--------------------|-------------------------|
| Balanced HHP | 111                    | 136                    | 178               | 98                | Yes                | 2284                    |
| HHP % / leg  | 21%                    | 26%                    | 34%               | 19%               |                    |                         |

**Figure 8 – ML-designed Node Split HHP data**

With good HHP balance and the ability to quickly double capacity via node segmentation this serving area will not require significant capacity upgrade activity for several years at nominal HSI bandwidth growth rates.

## 5.2. Implementing and Scaling Node Split Machine Learning Environment

The benefits of implementing a Machine Learning Node Split Design Environment are numerous and give the operators the analysis capability that would be cost-prohibitive in a standard design environment. For example, the nodes that are currently on the high-contention report could be quickly run through the ML design tool, under different business rules, and a Capital Expense projection could be created to guide

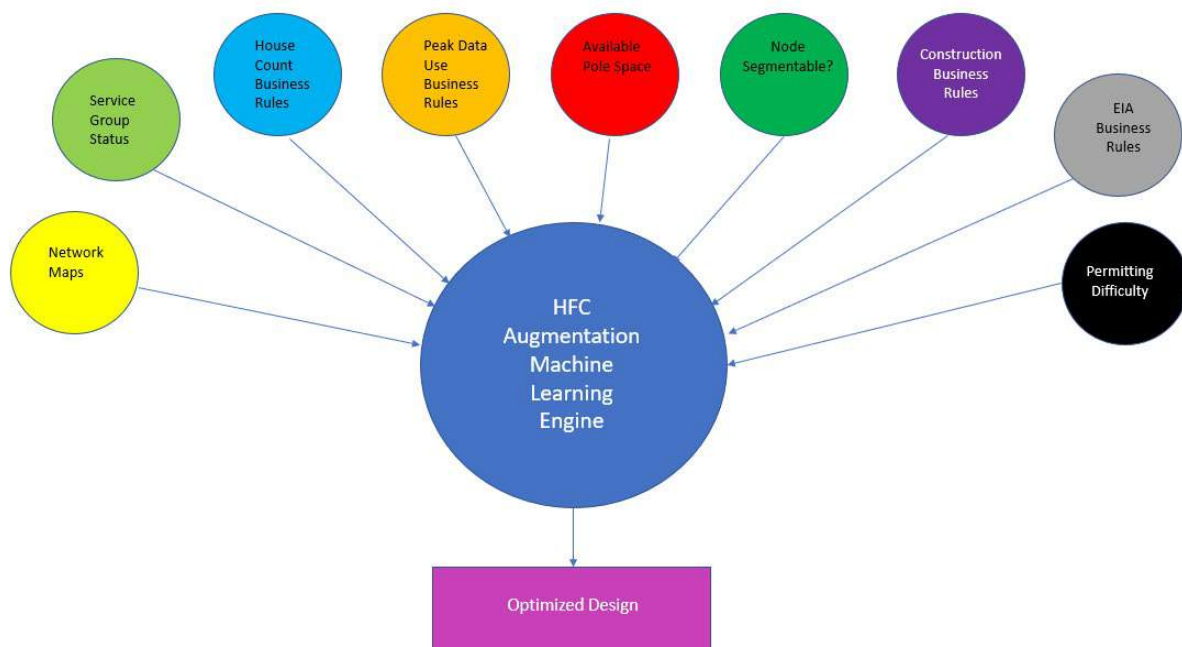
the budgeting process, and to inform the business about the tradeoffs involved where minimal construction augments result in more frequent network touches.

One challenge in implementing is the ML platform requires local data access – if your network is using a centralized network map platform such as GE SmallWorld or Synchronoss SpatialNet the maps would need to be “checked out” for data extract to the ML environment.

## 6. Creating a Holistic Network Capacity Augmentation Environment

As Machine Learning is integrated into the Network Planning and Operations process and platforms, as more systems and data are connected much more comprehensive design and analysis could be accomplished.

The chart in Figure 9 below shows some example inputs that if connected to a ML Prediction Engine could provide the network operator with planning information that could be calculated rapidly on a per-node basis for the entire network without significant human effort.



**Figure 9 – Possible inputs to holistic HFC Analysis, Design, and Planning Engine**

One possibility would be to connect the HSI network capacity platform and provide the ML environment data consumption information at a Peak-use-by-household basis. The ML business rules could be modified to optimize the node split based on real data usage versus just simply dividing the homes passed. A further business rule could identify the top 1% users and design and create BOMs for fiber connections to them to remove their data consumption from the HFC network – and cost scenarios could be quickly compared.



## 7. Conclusion

Network planning via Machine Learning is in its infancy but holds exciting promise to provide accurate and rapid automation of time-consuming network design activities. Network Operators should consider implementing network-based Machine Learning and over time connecting all network management platforms to enable comprehensive network design and capital expense planning.

## Abbreviations

|      |                                    |
|------|------------------------------------|
| AI   | Artificial Intelligence            |
| AR   | Augmented Reality                  |
| BOM  | Bill of Materials                  |
| CAGR | Composite Annual Growth Rate       |
| DS   | Downstream                         |
| EIA  | Electronic Industries Association  |
| GBM  | Gradient Boosting Machine          |
| HFC  | Hybrid Fiber-Coax                  |
| HHP  | HouseHolds Passed                  |
| HSI  | High Speed Internet                |
| IDE  | Integrated Development Environment |
| ML   | Machine Learning                   |
| MSO  | Multiple System Operator           |
| US   | Upstream                           |
| VR   | Virtual Reality                    |

## Bibliography & References

Cloud Computing: Differences Between the AI and ML. <https://www.comparethecloud.net/articles/cloud-computing-differences-between-the-ai-and-ml/>

Machine Learning: An In-Depth Guide - Overview, Goals, Learning Types, and Algorithms  
<https://www.innoarchitech.com/blog/machine-learning-an-in-depth-non-technical-guide>

# **Bringing Service Visibility Into The Light With CPRI As A Service**

A Technical Paper prepared for SCTE•ISBE by

**Bill Beesley**

MSO Business Development Manager  
Fujitsu  
2801 Telecom Parkway, Richardson, TX 75098  
972.479.2098  
bill.beesley@fujitsu.com

# Table of Contents

| <b>Title</b>                                | <b>Page Number</b> |
|---------------------------------------------|--------------------|
| 1. Introduction.....                        | 3                  |
| 2. Are we stuck in the dark?.....           | 3                  |
| 3. New technologies to change the game..... | 4                  |
| 4. Conclusion.....                          | 7                  |
| Abbreviations.....                          | 8                  |
| Bibliography & References .....             | 8                  |

## List of Figures

| <b>Title</b>                              | <b>Page Number</b> |
|-------------------------------------------|--------------------|
| Figure 1 - Unified Software Control ..... | 6                  |

## 1. Introduction

The Common Public Radio Interface (CPRI), which is commonly used to connect cell sites and base stations, imposes ultra-low latency and transparent communications requirements that effectively force operators to provide dark fiber to carriers who purchase their services. But using dark fiber for a single service to an end customer is not only a wasteful use of a high-value asset, it is operationally inefficient. Dark fiber gives no visibility into the services the customer is running, and consequently the operator has no proactive way to monitor and manage the health of these services. Often, operators are alerted to outages when the customer calls to report them, which is, at best, a weak mechanism for support visibility. This paper will outline how new technologies—such as low latency Ethernet and hybrid active/passive networking—enable operators to create CPRI-as-a-Service offering that overcomes this key disadvantage, improves overall management and maintenance of their xHaul networks, and lets them reclaim valuable unused fiber assets.

## 2. Are we stuck in the dark?

One of the little known secrets in the industry is that Cable Company owned fiber services the vast majority of cell towers in the United States. Pat Esser, Chief Executive Officer for Cox Communications shared during a Fox Business interview from the floor of the Consumer Electronics Show in 2019 that they provide the fiber to 82% of the cell towers that are located in their footprint.<sup>1</sup> Within the United States the fortunate convergence of a regulatory environment that created a robust private cell tower ecosystem and that resulted in the MSO community investing a rich fiber deep network infrastructure resulted in the majority of wireless carriers opting to lease xHaul services from the cable companies rather than investing in this infrastructure themselves. In the early days of xHaul, SONET was often used to provide traditional TDM based transport of bandwidth from the cell tower back to the wireless service provider's telecommunications facilities. Over time as higher performance and lower cost Ethernet transport became available, this became the preferred transport mechanism commonly served via MPLS services from the cable company. In parallel to this transport evolution the wireless telecommunications equipment located at the towers also evolved an important technology. The Common Public Radio Interface specification or CPRI, created a mechanism which allowed the Baseband Unit (BBU) that performs all the signal processing and management to be physically separated from the wireless radio, allowing Remote Radio Units (RRUs) to be located several thousand feet away. The original intent of the CPRI specification was to allow distributed RRUs and their antennas to be remotely located on the tower with the baseband processing systems located in buildings a short distance away. Because of this short distance expectation, the CPRI specification has very strict network timing and low latency requirements. The specification also requires transparent layer one connectivity to in-band communication between the BBU and RRU. With an allowed latency of 75 microseconds, wireless operators realized that they could take advantage of economies of scale and lowered operations efforts by centralizing their BBUs in facilities located up to 20 kilometers away from the cell towers. This low latency, transparent communication link between the radios and base stations is commonly referred to as “fronthaul” and unlike legacy backhaul networks could not be supported via MPLS or Ethernet service offerings from the MSO forcing them to use dark-fiber to support the transport needs of the wireless carriers.

But dark fiber is an extremely valuable asset. In the United States the average cost to build out fiber infrastructure is around \$100,000 per mile, assuming a blended aerial and underground mix. Not only is that a difficult return on investment model, especially considering that the perceived service value of the carrier purchasing dark fiber service is that it is in the range of lit Ethernet services, repaying out a fiber build in the tens of thousands of dollars from a single cell tower is a real challenge. Now it should be pointed out here that when fiber is built, we do not just pull a pair of fibers as the general practice is to pull a sheath of fibers with hundreds of fibers. While these extra fibers can be used to facilitate other

network needs such as multiservice distributed access or other commercial customer buildouts, the needs of the other services may not be optimally placed in relation to the cell tower. The question should also be asked, “even if I can pay the cost of the build and/or make some profit off providing dark fiber service....is that the best fiscal use of this high value asset?” A single fiber strand can provide services to hundreds, and sometime thousands of residential and commercial subscribers. Dark fiber service delivered to the tower orphans that asset to a single customer.

While the economics of providing dark fiber as a service are arguable, there will always be a need for dark fiber and there will always be companies willing to offer it, what is not debatable is that dark fiber is difficult to monitor for service health.

Wireless carriers are notorious for their high expectations of service health and reliability. Since they are service providers themselves, they have an accountability to what can be hundreds of their customers who may be impacted by outages. Since the MSO will have no visibility into the network equipment connected to the dark fiber, they will also have not be able to monitor service health. This often means that the first notification of a customer affecting outage the MSO will likely receive is when the customer calls to tell them they think they have an issue with their fiber. Since the MSO has no visibility into the services running on the fiber they are left blind to the root cause of service degradation. Is the cause an impairment caused by damage to the fiber, an issue at the CPRI or Ethernet layer, or some other issue? Without visibility into the service layer the opportunity for finger pointing between customer and provider in the fog of war that comes with service degradation effecting multiple downstream customers becomes a common event that can negatively impact the relationship between MSO and carrier.

So, is the option only to offer dark fiber with its high cost and low service value because the needs of the CPRI protocols cannot be support by traditional Ethernet or MPLS? The good news is that technology has evolved and the options are no longer just a binary dark fiber or nothing to satisfy the requirements of wireless carriers.

### **3. New technologies to change the game**


One set of technologies that is a game-changer in the ability to introduce concepts known as low latency Ethernet and/or Time Sensitive Networking. Time Sensitive Networking is a series of standards built on the 802.1 working group standards that allow Ethernet to support very low latency and jitter as well as introducing strict network time synchronization required to so support CPRI fronthaul applications as a service.<sup>2</sup> The technology advancements introduced by TSN have allowed for the introduction of CPRI as a service. Much like traditional SONET, ATM and Ethernet services, CPRI services allow the MSO to deliver the service requirements of ultra-low latency and strict network time synchronization while also allowing them to manage, monitor and maintain the quality of the service to ensure a consistent and high customer experience level. Still, having the technical capabilities to deliver a CPRI service may not be sufficient by itself to wean carriers off their dark fiber addiction. Since there are fiber over builders often willing to construct in fiber, many MSOs have been reluctant to push their customers to a CPRI service offering out of concern that another fiber provider will overbuild their footprint, a very real risk. So, we need to discuss how to increase the customer perceived value of the service and discuss new service capabilities that will improve the customer experience.

But first, we need to touch on one other networking technology that is also viable for carrying carrier distributed antenna network, passive DWDM. It is quite common for larger wireless carriers who also have some of their own fiber to use passive DWDM with tunable optics to carry distributed antenna traffic in their networks. There are a good variety of 10G and 25G single-and dual-fiber options that can plug directly into the remote radio units to carry the traffic back to the baseband unit for processing.

Much like DWDM is used for DAA and business services traffic in the MSO network, this is an excellent low cost option for metro networks. However, just as with dark fiber, passive DWDM lacks service layer visibility and monitoring. Unless the MSO has visibility into the edge equipment carried on the passive network, they will have no idea that there is service degradation until the customer calls them and even then, it will be just as difficult as with dark fiber to quickly troubleshoot and isolate the root cause.

As stated previously, wireless carriers justifiably, have a high expectation for service availability and quality. What is needed to elevate the level of these technologies to a true service offering is a way to make the network more intelligent, more automated and easier to manage and maintain, and more transparent to both the MSO and to the wireless carrier.

It is becoming increasingly more common for modern network transport and access equipment to be equipped with standardized APIs and YANG models to allow them to be part of a Software Defined Networking or SDN ecosystem. SDN allows operators to evolve the paradigm of network management from provisioning to programming. By imagining a network as a programmable set of components that can obey defined rule sets and consistently report information such as network topology and system and service health, we can take advantage of SDN technologies to create new service capabilities.



**Figure 1 - Unified Software Control**

In figure 1 we see a block diagram example of how we can use software control over the end to end network to provide a unified operations view for both the MSO and the wireless operator. Taking advantage of open APIs we have a transport SDN controller that is within the MSO network and a wireless SDN controller in the wireless operator's network. Both the wireless and transport networks are controlled by an orchestration system and service operations that both operator and customer have delegated views into but with walled garden control. Both operator and customer can "see" the status of the other network and their orchestration systems can respond to and honor network demands from the controllers as well as provide unified topology and management views of the end-to-end network and its health.

This mixed management methodology will operate and behave based on service rules agreed to between operator and customer. Operations engineers familiar with traditional network management systems and more importantly the lack of multi-vendor and multi-layer management capabilities will undoubtedly be skeptical of the image in figure 1. This has traditionally been true as both network equipment and management systems have been closed, monolithic ecosystems lacking a high degree of flexibility and extensibility. But, modern network automation software is capable of introducing this new functionality into existing operations and control systems. Today we can take advantage of open APIs and microapps to introduce functionality into operations systems incrementally without the need to tear out the existing infrastructure. This method is also capable of connecting multiple vendor systems because it can integrate either with existing management platform northbound APIs or integrate directly with the network infrastructure itself via document YANG model. Both of these mean that no longer must a new, monolithic system be stood up next to existing infrastructure. The MSO can add the new functionality a service at a time using microapplications and microservices.

One example of where using software automation to modernize networks is improving the functionality of passive DWDM networks. As previously stated, traditional passive DWDM networks allowed for very

limited visibility into the topology and health of the underlying transport infrastructure because there was no direct way to query this equipment. This meant, as an example, If an operator wanted to provide a service level topology view of their network their technicians would be required to manually provision wavelength services as they were added, moved, or deleted. Since manual efforts like this are both labor intensive and fraught with potential for error, the result was that not many if any operators have a live topology view of services running on their passive optical networks for monitoring and troubleshooting purposes. In addition to a lack of service visibility as passive networks require the client pluggable to be correctly tuned to the wavelength supported by the port into which they are plugged, they increase the operational effort to provision network additions or changes. They also increase the potential for error as the technician can select the wrong port or pluggable.

Modern smart pluggables overcome this issue as they are tunable across the full optical spectrum and are able to determine which wavelength they should auto-tune to in order to establish service. These pluggables are also capable of some service monitoring to determine the health of the optical link. By using software intelligence to monitor this, a micro-application can be introduced into the network management system to determine which wavelength the pluggable is tuned, correlate this with the known passive topology, and then draw a visual representation of the service path through the passive network. This can then be later used for system alarming, showing the network operations team that not only that the pluggable is in alarm, but the expected path through the network. This provides additional service context that can be useful in quickly isolating and resolving service faults. Additionally, the operator can provide the smart pluggables to customers who are connecting in their remote radios or other devices to the passive network. Since the pluggable will auto-provision and tune to the appropriate wavelength, the wireless operator can self-install their equipment once fiber has been terminated at the appropriate demarcation point. As the pluggable can be monitored by the MSO, they now have service visibility for the optical layer without having to have access to the customer's edge equipment. This topology is known as Hybrid Active/Passive and can greatly improve the customer experience while improving MSO installation and operations efforts.

## 4. Conclusion

In addition to allowing for multi-layer, multi-vendor operations, these modern systems often have well documented APIs it becomes practicable to implement intersystem communications and control between operator and customer. Since the implementation can be done on a per-function application via a micro-app or micro-service, this integration work can be done within the context of a specific customer agreement and the cost of implementation can be tied to the revenue generated by that agreement. Open systems and open architectures in the transport and wireless networks and their control systems mean we can now extend automation and visibility of the network to allow for lower operations effort and provide a better customer experience.

Advances in time sensitive and hybrid active/passive networks now allow the capability to carry wireless operators xHaul traffic in an efficient and reliable manner negating the use of dark fiber as the only method to support the needs of wireless operators. These technologies are deliver a better customer experience by enhancing the MSO's visibility into the service health, which is not feasible with dark fiber. These technology advances are likely enough in their own right for MSOs to justify offering CPRI as a service in lieu of dark fiber.

But the industry should not stop there when considering their xHaul service offerings. With modern software intelligence, further improvements can be added to the service that increase collaboration and trust between the MSO as a transport operator and the wireless carrier. Ultimately these will lead to an increase in customer satisfaction and increase the product "stickiness".



## Abbreviations

|      |                                               |
|------|-----------------------------------------------|
| DWDM | dense wave division multiplexing              |
| DAA  | distributed access architecture               |
| MSO  | multiple system operator                      |
| HFC  | hybrid fiber-coax                             |
| HD   | high definition                               |
| Hz   | hertz                                         |
| ISBE | International Society of Broadband Experts    |
| SCTE | Society of Cable Telecommunications Engineers |

## Bibliography & References

Fox Business interview of Pat Esser from the Consumer Electronics Show 2019 by Liz Claymon  
<https://www.facebook.com/LizClaman/videos/518267898681406>

Time Sensitive Networking Task Group <https://1.ieee802.org/tsn/>

# **Make the Most of What You've Got: How Cable Modems Can Deliver Economical Cell Site Transport**

A Technical Paper prepared for SCTE•ISBE by

**Bill Beesley**

MSO Business Development Manager

Fujitsu

2801 Telecom Parkway, Richardson, TX 75098

972.479.2098

[bill.beesley@fujitsu.com](mailto:bill.beesley@fujitsu.com)

# Table of Contents

| <b>Title</b>                          | <b>Page Number</b> |
|---------------------------------------|--------------------|
| 1. Introduction.....                  | 3                  |
| 2. A Convergence of Technologies..... | 3                  |
| 3. Choices, Choices.....              | 4                  |
| 4. An evolution.....                  | 5                  |
| 5. Conclusion.....                    | 5                  |
| Abbreviations.....                    | 5                  |

## 1. Introduction

MSOs seeking to increase service delivery speed and reduce the high cost of cell site transport over fiber now have a new option. Low-latency DOCSIS has made the HFC network viable for midhaul and backhaul transport applications. The Telecom Infra Project's vRAN Fronthaul Project Group has developed a means to use cable modems for some cell site transport applications. Using SDN to provide unified management and control of the fiber xHaul; DOCSIS BSOD services; and the RAN, operators can now create the consistent operational views and control needed to fully support this new application for cable modems. This paper will outline the future of wireless transport and discuss how to operationalize cable modems into the existing fiber transport architecture. Topics will include how to implement automated setup and provisioning of network demands, along with a single view of service health, in an overall consistent provider experience whether they use traditional fiber transport or cable modems for xHaul service delivery.

## 2. A Convergence of Technologies

The Common Public Radio Interface specification or CPRI created a mechanism that allowed the Base Band Unit (BBU) that performs all the signal processing and management to be physically separated from the wireless radio allowing Remote Radio Units (RRUs) to be located several thousand feet away. The original intent of the CPRI specification was to allow distributed RRUs and their antennas to be remotely located on the tower with the baseband processing systems located in buildings a short distance away. Because of this short distance expectation, the CPRI specification has very strict network timing and low latency requirements. The specification also requires transparent layer one connectivity to in-band communication between the BBU and RRU. With an allowed latency of 75 microseconds, wireless operators realized that they could take advantage of economies of scale and lowered operations efforts by centralizing their BBUs in facilities located up to 20 kilometers away from the cell towers. The industry uses the term "fronthaul" to describe the access layer connection between the centralized base band units to the distributed antenna system. Because of fronthaul's low latency requirements, optical fiber based solutions are the common method for building this access layer. But with millions of distributed antennas predicted to be deployed to support continued growth of the 5G era, especially in remote and rural areas, the industry is seeking to advance other technologies like cable modems to be able to support fronthaul requirements.

The Telecom Infrastructure Project vRAN Fronthaul Project Group has set out to develop solutions that would allow the connection between the baseband unit and the remote radio unit to have much lower latency requirements than currently allowed by the CPRI specification. The first round of technology labs testing access technologies such as PON, G.Fast and cable modems was completed in 2019. CableLabs® served as the community lab that provided proof of concept testing for the use of cable modems and they were able to successfully demonstrate support for fronthaul services over an HFC network in their lab environment. While non-ideal fronthaul is still in the proof of concept stages as of the time of this paper, the results are promising and could be standardized and supported by the radio vendor products in the next few years. This would do much to accelerate the deployment of 5G services and give cable operators another viable choice for access network transport of fronthaul. (see <https://telecominfraproject.com/vran/>)

In parallel to this activity, CableLabs® has developed mechanisms to significantly improve the latency in HFC networks. Because of the communications handshakes between the cable modem and the CMTS and the nature of the way DOCSIS schedules upstream timeslots, latency in the HFC network has generally been way too high to ever consider it for use in the split RAN architectures described above. The industry needed to be able to support applications like gaming and split RAN xHaul which have very low latency demands. CableLabs® accomplished this by introducing low latency DOCSIS for gaming and low latency xHaul for wireless transport. Many papers go into the details of how both of these work but a simplified description is that by modifying the scheduling mechanism to be preemptive and allow for parallelism it becomes possible to get the latency of the HFC network down into the one to three millisecond range. While this does not currently support fronthaul, in combination with non-ideal fronthaul at some point in the near future it will be entirely possible to support all xHaul architectures via the HFC network.

But as we move into the “can we” question of using HFC for fronthaul, we also need to consider the “should we”. No network tool is universal and there are several operational considerations for using coax versus fiber to support the deployment of fronthaul based radio units

### **3. Choices, Choices**

The first consideration is the utilization of the asset itself. Wireless deployments are not “bursty” like residential or other commercial data users. Today, a radio unit will use the maximum amount of bandwidth with only one radio connected to it. In addition, the preemptive nature of low latency xhaul means that some amount of upstream bandwidth will need to be allocated in case the radio needs to communicate with the base band unit. When planning for future cable modem based fronthaul deployments, the question will need to be asked if the HFC bandwidth could be used for other revenue generating services. In addition, HFC network reliability is significantly lower since all but one MSO does not use a ring architecture. Depending on the sizing of the wireless service group, thoughtful consideration will need to be given as to whether it may make economic and operational sense to deploy fiber even where coax may be readily available. Of course, where fiber assets are unavailable, or too expensive to extend, coax will still be a great option for fronthaul. Operations staff will just need to measure the options.

An old telecom service adage goes something like, “it is easy to deploy the first service. It is deploying the next ten thousand that is hard.” As the paradigm for fronthaul services moves into the HFC world, operations staff will need to be retooled and retrained to support a new class of customer. Wireless carriers have much higher service level expectations than existing residential and commercial HFC based customers. This, along with the introduction of new technologies such as remote radio units will increase expectations placed on frontline operations staff. One example of this change was experienced a bit over ten years ago when the first cable modem based, strand mount wireless hotspots were introduced into the market. A major cable operator in the US decided to deploy this new technology to support wireless data based services in public spaces and event areas. These devices were mounted on the existing HFC strand and obtained both their access and power from the coax to deliver a public hot spot service. What the cable operator quickly learned was that the outside plant technicians didn’t have a working knowledge of cable modems and wireless access points and that the wireless data engineers did not have experience working with HFC hardline, plant power and even if they did, they didn’t have bucket trucks. It became quickly clear that these groups were going to have to develop new tools, training and working methods in order to successfully ensure the delivery of the new service.

When considering adding coax as a service delivery mechanism for fronthaul, operations teams need to plan on how they will deploy the new service in a consistent and repeatable manner. How will service

adds, moves and changes be managed quickly and efficiently? Who will monitor and manage service health and how will they accomplish this? There will be new tools, new methods and of course, new training necessary. Frontline operations leadership needs to carefully document and prepare their teams for the changes that will come.

Another consideration often overlooked is that there are a large number of wireless carriers currently supported by the MSO fiber infrastructure. In the US, the cable companies own the majority of fiber infrastructure delivering services to wireless carriers. Pat Esser, Chief Executive Officer for Cox Communications shared during a Fox Business interview from the floor of the Consumer Electronics Show in 2019 that they provide the fiber to 82% of the cell towers that are located in their footprint. Due to the fiber rich nature of the HFC architecture, it is likely other MSOs are enjoying similar service numbers within their footprint. As the cable industry evolves toward delivering some wireless xHaul via cable modems there will need to be a mechanism to provide unified and consistent management and support for the customer where their services are running on both fiber and coax networks.

## 4. An evolution

The industry will need to evolve towards support for a unified topology paradigm that allows the end-to-end management both fiber and coax networks along with the wireless radio network. New technologies like software defined networking controllers will be required to allow for collaboration across these network domains in order to not only support functionality, such as 5G network slicing, but also to allow operations teams to be effective. Future networks will require operations teams to shift from managing network equipment to managing customer services. Services will need to be automatically created at both the fiber access and DOCSIS network based on service demands from the wireless access layer and then effectively chained together to provide an end-to-end service. To avoid the need to add additional operations resources as new, more complex services continue to be added to the network, it will become the software management layer that will be the critical linchpin in the operators network. Software interfaces will need to become more open. Systems will need to allow for more intelligence and automation. Emerging technologies such as AI/ML will need to be introduced into the management of the end-to-end network. The alternative is operational overload.

## 5. Conclusion

The cable industry is at another turning point in service delivery. Wireless deployments will continue to grow exponentially to meet growing demand for more bandwidth and lower latency. As the deployment of fiber deep assets to support HFC allowed them to secure the access transport for the majority of the cell towers deployed in the past, the MSO community is well positioned to support the next round of wireless growth. But they will need to prepare more than their network technologies to support this next wave of new services. They must also invest in and prepare their operations systems and teams. By doing so, cable networks will continue to be the dominant method for delivering access transport for wireless deployments

## Abbreviations

|       |                                          |
|-------|------------------------------------------|
| AI/ML | artificial intelligence/machine learning |
| BSOD  | business services over DOCSIS            |
| BBU   | base band unit                           |
| CMTS  | Cable modem termination system           |

|        |                                                 |
|--------|-------------------------------------------------|
| CPRI   | common public radio interface                   |
| DOCSIS | data over cable service interface specification |
| HFC    | hybrid fiber-coax                               |
| G.fast | ITU-T G fast access to subscriber terminals     |
| MSO    | Multiple system operator                        |
| PON    | Passive optical network                         |
| RAN    | radio access network                            |
| RRU    | Remote radio unit                               |
| SDN    | Software defined networking                     |
| vRAN   | virtualized radio access network                |
| xHaul  | fronthaul, mid-haul, backhaul                   |

# **With 1.2GHz of spectrum are we moving to a channelized per room architecture for the home – enabled by Wi-Fi 7**

A Technical Paper prepared for SCTE•ISBE by

**Charles Cheevers**

CTO Home Networks Solutions  
CommScope  
Charles.Cheevers@commscope.com

**Ian Wheelock**

Engineer Fellow  
CommScope  
Ian.Wheelock@commscope.com

**Kurt Lumbatis**

Distinguished Software Engineer  
CommScope

Kurt.Lumbatis@commscope.com

**Kamal Koshy**, Charter Communications

**Ahmed Bencheikh**, Charter Communications



# Table of Contents

| Title                                                                               | Page Number |
|-------------------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                                | 4           |
| 2. Todays home – legacy g/n and ac devices and emerging Wi-Fi 6 .....               | 7           |
| 3. Service Providers and the introduction of Wi-Fi 6.....                           | 11          |
| 4. Setting the stage for Wi-Fi 7 entrance – Wi-Fi 6E .....                          | 16          |
| 5. Low Power Indoor Mode to AFC controller Standard Power.....                      | 23          |
| 6. Wi-Fi 7 and the 320MHz channel what does it mean for the home.....               | 25          |
| 7. What do we do with 2.4GHz and 5GHz spectrum in the Wi-Fi 7 era.....              | 33          |
| 8. Challenges and opportunities for Wi-Fi ecosystem .....                           | 35          |
| 8.1. Enhanced quality of service for applications like VOIP, voice and gaming ..... | 35          |
| 8.2. Onboarding of IoT devices in Wi-Fi networks .....                              | 36          |
| 8.3. Virtualization in Wi-Fi networks.....                                          | 36          |
| 8.4. Operational improvements and network monitoring.....                           | 36          |
| 8.5. Devices supporting multiple networks .....                                     | 36          |
| 9. Conclusion .....                                                                 | 37          |
| Abbreviations.....                                                                  | 38          |
| Bibliography & References .....                                                     | 38          |

## List of Figures

| Title                                                                                       | Page Number |
|---------------------------------------------------------------------------------------------|-------------|
| Figure 1 - The Wi-Fi Standards.....                                                         | 4           |
| Figure 2 - The evolution of Wi-Fi from 5 to 7 .....                                         | 5           |
| Figure 3 - Potential overlapping Wi-Fi DOCSIS GW solutions in the 2020-2025+ timeframe..... | 6           |
| Figure 4 - The Wi-Fi Eras and where we are going.....                                       | 6           |
| Figure 5 - Wi-Fi Speed evolution.....                                                       | 7           |
| Figure 6 - Average time to replace Smartphone in the US .....                               | 8           |
| Figure 7 - Todays Wi-Fi architectures .....                                                 | 9           |
| Figure 8 - Dual Band Concurrent AP (2.4+5GHz).....                                          | 10          |
| Figure 9 - Tri Band Concurrent (2.4+5GHz Low+5GHz High) .....                               | 11          |
| Figure 10 - The Primary Features of Wi-Fi 6 .....                                           | 12          |
| Figure 11 - 20MHz RU allocation mapping .....                                               | 12          |
| Figure 12 - 40MHz RU allocation mapping .....                                               | 12          |
| Figure 13 - DL MU-MIMO + OFDMA.....                                                         | 14          |
| Figure 14 The decision points for Wi-Fi 6 introduction .....                                | 14          |
| Figure 15 - Wi-Fi 5 vs Wi-Fi 6 2x2 Client Performance Capability (Mbps) .....               | 15          |
| Figure 16 - STB client performance showing 2x2 Wi-Fi 6 outperforming AC 4x4.....            | 15          |
| Figure 17 - STB client performance showing 2x2 Wi-Fi 6 outperforming AC 4x4.....            | 16          |
| Figure 18 - US 6GHz spectrum and channel allocation to Wi-Fi .....                          | 17          |
| Figure 19 - Channel Availability over Spectrum.....                                         | 17          |
| Figure 20 - TBC 2.4+5GHz AP.....                                                            | 18          |
| Figure 21 - 4x4 to 4x4 Wi-Fi 6E performance at 5dBm/PSD - LPI specification.....            | 19          |

|                                                                                                                                     |    |
|-------------------------------------------------------------------------------------------------------------------------------------|----|
| Figure 22 - Switchable 5/6 designs will bring in cost and complexity levels that may not be acceptable to Operator cost points..... | 20 |
| Figure 23 - ARRIS VAP2400 Wi-Fi Video Bridge.....                                                                                   | 21 |
| Figure 24 - External Wi-Fi 6E dongle for existing GW.....                                                                           | 21 |
| Figure 25 - Can 2.4GHz be separated from the Wi-Fi mesh 5/6/7 devices in some future architecture....                               | 22 |
| Figure 26 - Add 6GHz to AP, but extraction all 2.4GHz Radios.....                                                                   | 23 |
| Figure 27 - Current potential channel numbering scheme for U-NII5 to U-NII8.....                                                    | 24 |
| Figure 28 - Immersive display on the LightField Video Looking Glass display at 8K levels.....                                       | 26 |
| Figure 29 - Basic single 30dBm - 8x8 or 4x4 Cross Home Wi-Fi 7 Architecture .....                                                   | 27 |
| Figure 30 - 8x8 AP/GW at 30dBm in both 320MHz and 160MHz channels .....                                                             | 27 |
| Figure 31 - 4x4 performance at 320MHz and 160MHz at 30dBm.....                                                                      | 28 |
| Figure 32 - 8x8 performance of 320MHz and 160MHz at 2W radiated Power .....                                                         | 29 |
| Figure 33 4x4 performance of 320Mhz and 160MHz at 2W radiated Power.....                                                            | 29 |
| Figure 34 Performance of VLPI in room potential low power solution.....                                                             | 30 |
| Figure 35 - Wi-Fi 7 backhaul - Wi-Fi 6E in room - Ultimate Home Wi-Fi ? .....                                                       | 31 |
| Figure 36 - use of Wi-Fi backbone and in room 4.7GBps networks.....                                                                 | 31 |
| Figure 37 - Wi-Fi 7 and Wi-Fi 6E - driving high capacity low latency services .....                                                 | 32 |
| Figure 38 - 2x2 320MHz Wi-Fi 7 vs 4x4 160MHz Wi-Fi 6E .....                                                                         | 33 |
| Figure 39 - OFDMA with 2MHz DCM .....                                                                                               | 34 |
| Figure 40 - Improved Wi-Fi 6 IoT co-existence with ZB/BT .....                                                                      | 34 |
| Figure 41 - Does 2.4GHz separate out from 5/6GHz in the home over time ?.....                                                       | 35 |
| Figure 42 - Potential synergy between different Operator Wireless networks.....                                                     | 37 |

## List of Tables

| Title | Page Number |
|-------|-------------|
|-------|-------------|

NO TABLE OF FIGURES ENTRIES FOUND.

# 1. Introduction

We are in a unique time in our history of Wi-Fi with three major changes happening in less than a 3 year period that fundamentally challenge us to look at our traditional view of Wi-Fi home architectures in the home. In particular how a Service Provider keeps pace with change and reaps the benefit of the increase in bandwidth, increase in coverage and decrease in latency while at the same time providing their customers with a clear and concise device and services roadmap.

The 3 fundamental changes are

- (i) Wi-Fi 6 – just starting to deploy this year for most service providers in their primary Gateway or Access Point. Wi-Fi 6 extender solutions typically lagging behind the primary AP function.
- (ii) Wi-Fi 6E and the addition of 6GHz spectrum. In particular in the US the allocation of 1.2GHz of new unlicensed spectrum has now fundamentally opened up a new approach to how Wi-Fi and all wireless services play out in the home. And provide a path to even higher bitrate immersive video services and lower latency services over Wi-Fi.
- (iii) IEEE802.11be and probably the next Wi-Fi 7 standard which takes the opportunity now of 6GHz spectrum and opens up single channel capabilities to support 320MHz channels from the present 160MHz channels supported today in Wi-Fi 6/6E.

| Date    | Standard | Wi-Fi | Frequency  | Bandwidth MHz    | MIMO  | Bitrate   |
|---------|----------|-------|------------|------------------|-------|-----------|
| 1997    | 802.11   |       | 2.4GHz     | 20               | NO    | 2Mbps     |
| 1999    | 802.11b  | 1     | 2.4GHz     | 20               | NO    | 11Mbps    |
| 1999    | 802.11a  | 2     | 5GHz       | 20               | NO    | 54Mbps    |
| 2003    | 802.11g  | 3     | 2.4GHz     | 20               | NO    | 54Mbps    |
| 2009    | 802.11n  | 4     | 2.4+5GHz   | 20/40            | 4x4   | 600Mbps   |
| 2013    | 802.11ac | 5     | 5GHz       | 20/40/80/160     | 8x8   | 6.9Gbps   |
| 2019    | 802.11ax | 6     | 2.4+5GHz   | 20/40/80/160     | 8x8   | 9.6Gbps   |
| 2021    | 802.11ax | 6E    | 2.4+5+6GHz | 20/40/80/160     | 8x8   | 9.6Gbps   |
| 2023/24 | 802.11be | 7     | 2.4+5+6GHz | 20/40/80/160/320 | 16x16 | 30-40Gbps |


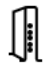


**Figure 1 - The Wi-Fi Standards**

While Wi-Fi took 20 years to finally revise its MAC for the creation of IEEE802.11ax and subsequent Wi-Fi 6 standard we now have both 6E and Wi-Fi 7 happening over a 3 year period so need to now look ahead to the next 3 years and how the

- Transition from current home solutions of mixed mode g/n and ac to emerging Wi-Fi 6 AP and the start of third party Wi-Fi 6 clients
- Introduction of specific Wi-Fi 6 clients for Service Providers – typically first STB and Extenders
- Planning for use of 6GHz spectrum and the introduction of Wi-Fi 6E. Current discussions are around potential introduction by all new Tri-band Gateway or AP or with the introduction of a Wi-Fi 6E extender that supports augmentation for Wi-Fi 6 as well as first Wi-Fi 6E usage.

- Use of Low Power Indoor power levels for Wi-Fi 6E and the utilization of higher Standard power levels utilizing Automated Frequency Control mechanisms
- Exploitation of the 33-66Gbps of Wi-Fi spectrum available with 6GHz channels alone for the next generation home and services
- Use of 320MHz high capacity channel for high capacity client applications or use for backhaul to 10GbE capability in the home.
- Separation of 2.4GHz and 5GHz usage from 6GHz channels for legacy solutions as well as potential new usages for these legacy spectrum channels. 2.4GHz for example is now in stark relief for its limited capacity versus the other frequency levels but also may have a key role to play in the home for long range 'Narrow Band' Wi-Fi IoT functionality. Its current role for being primarily a reliable long range coverage channel when 5GHz is not usable or the device types are cheap and only support 2.4GHz radios.

This paper will discuss these areas in more detail in the following sections and hopes to offer a good discussion on the parameters and decision points to consider in the next 3 years of fast moving Wi-Fi technology changes. It should help the Cable Operator with the transition decision points of moving from Wi-Fi 5 through 6 and 6E as well as Wi-Fi 7.

|         |                                                                                     | xBC              | Speeds      | Main Features                                             | Comment                                                                                                                                                                                                                                                                   |
|---------|-------------------------------------------------------------------------------------|------------------|-------------|-----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Mature  |    | Dual             | 1Gbps       |                                                           | 160MHz channels were supported – but clients did not typically support                                                                                                                                                                                                    |
| Current |   | Dual and Triband | 4x4 2.5Gbps | 4x speed<br>Lower latency<br>OFDMA<br>Better Battery Life | Wi-Fi 6 AP performs better for mixed Wi-Fi 4,5,6 clients than a Wi-Fi 5 AP.<br>Typically <b>25% improvement</b> in mixed Wi-Fi client environment<br>160MHz channels supported<br>Matched with 2.5GbE Ethernet support                                                    |
| 2021/22 |  | Triband          | 4x4 4Gbps   | Immediate use of Wi-Fi 6 features for the client          | Wi-Fi 6E allows ONLY Wi-Fi 6 devices with 6GHz support to operate in the U-NII5 6GHz band.<br><b>Immediate</b> full use of the performance of Wi-Fi 6 with 160MHz channels and QAM1024 modulation<br>Perfect solution for congestion (MDU) and low latency Wi-Fi services |
| 2023    |  | Triband          | 4x4 8Gbps   | 320MHz channels                                           | Wi-Fi 7 (currently IEEE802.11be specification) is extending the channel width to 320MHz. Potential for 10Gbps like <u>wireless backbone</u> across the home                                                                                                               |

**Figure 2 - The evolution of Wi-Fi from 5 to 7**

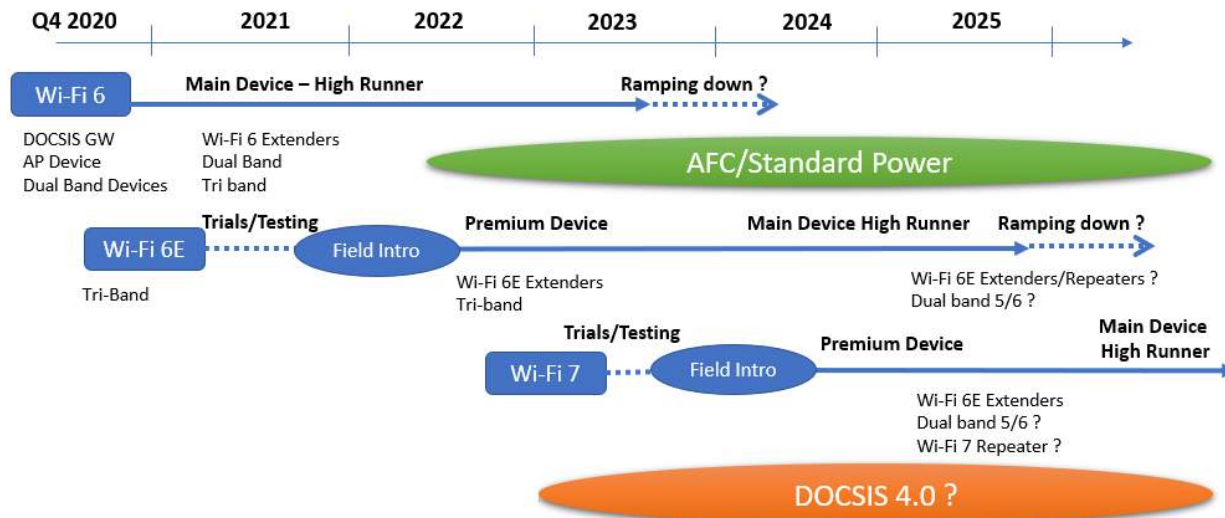
So, from an availability for a Cable Operator to be able to deploy new Wi-Fi technology today we are in the Wi-Fi 6 era with Cable Operators expected to deploy Wi-Fi 6E solutions end of 2021 or early 2022. This is then followed in 2023 by the ability to upshift once more with the proposed availability of Wi-Fi 7 solutions in market in 2023 – potentially even trialing as early as 2022.

Cable Operators need to invest in their device strategy wisely and ensure as usual that

- They minimize SKU's for keep down inventory and operational cost overheads
- Return the investment in their investment in GW and Extenders deployed
- Not confuse their customers with too many Wi-Fi changes
- Associate Wi-Fi changes with other consumer friendly features and performance metrics that they understand
- Align their Wi-Fi 6, Wi-Fi 6E and Wi-Fi 7 upgrades with their DOCSIS 3.1 to DOCSIS 4.0 evolution and increasing PON speed solutions.

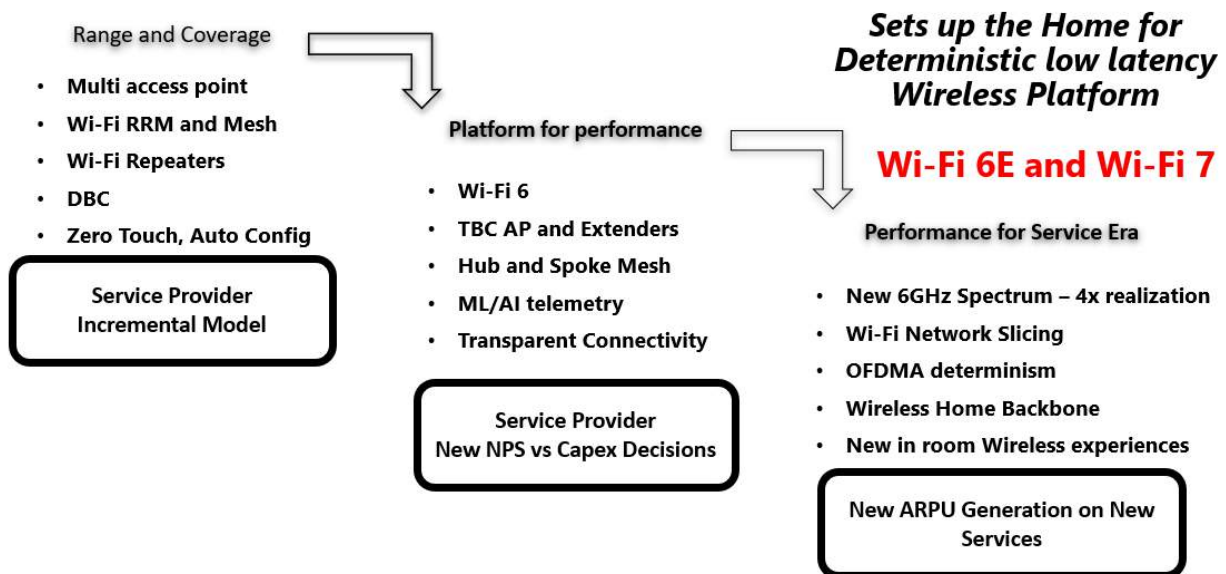
See below a potential evolution roadmap that highlights the introduction of both DOCSIS GW and Extenders to bring in the new Wi-Fi features. One potential alignment of Wi-Fi 7 now is DOCSIS 4.0 and these two technologies could be aligned for common introduction based on their current schedules. It

highlights the probability of Wi-Fi 6 GW today lasting well into the introduction of 6E and the continued use of Wi-Fi 6E devices in the Wi-Fi 7 domain particularly. It also brings up questions about Extension of Wi-Fi which we will also touch on more in the paper.



**Figure 3 - Potential overlapping Wi-Fi DOCSIS GW solutions in the 2020-2025+ timeframe**

Before we continue, we also want to outline some of the reasons why we need to make this investment in



**Figure 4 - The Wi-Fi Eras and where we are going**

We have been focused over the last 20 years of Wi-Fi on

- 802.11 b/g - getting untethered from the ethernet cable
- 802.11 n – using 5GHz spectrum which formed the basis of being able to support streaming IP video
- 802.11ac – which got us to 1Gbps speeds on Wi-Fi

- 802.11ax – which redefined the MAC and added the foundation for added Spectrum for Wider Channels

We got to a point where we also

- Initially had one AP per home
- Added Repeaters for ‘barely their capacity coverage’
- Improved with Wi-Fi Meshing solutions for better coverage
- Introduced new Wi-Fi 6 MAC and improved FEM technologies to be able to get more efficiencies in Mixed b/g/n/ac and ax client networks

But Wi-Fi still had difficulties getting 80MHz and 160MHz channels reliably enough, had high congestion issues in MDU’s and could never really get the determinism to feel reliable to have it as a foundation for new wireless high capacity and low latency services in the home.

Now with 6GHz spectrum (1.2GHz in the US) we have the new Era of Wi-Fi – Deterministic, Low latency high capacity platform – to build new Wi-Fi home services.

## 2. Today's home – legacy g/n and ac devices and emerging Wi-Fi 6

If we could change everyone’s Access Point to a Wi-Fi 6 device and every client to a Wi-Fi 6 device – we could take advantage of the improvements in features Wi-Fi 6 immediately. However, that is not possible for the Service Provider and even for the retail consumer. So, we currently have in everyone’s home a mixed environment for Wi-Fi where there are variances in the primary Access Point

- 802.11n and 802.11ac access points still dominate the distribution for homes
- Over 90% of deployed Access Points are Dual Band Concurrent
- The majority of Wi-Fi extenders are really repeaters with the majority of those still running dual band 802.11n
- The breakdown of client devices (taken from various CommScope sources) are spread across
  - o Pre AC (Wi-Fi 4) devices are > 43% of the current home devices (but in some regions of the world are as high as 71%)
  - o AC (Wi-Fi 5) devices are about 56% of the current home devices (in some regions of the world are as low as 27%)
  - o Wi-Fi 6 devices still only number about 1% of the current home devices




Figure 5 - Wi-Fi Speed evolution

Within the client base there are also new devices like the popular Wyze cameras that only support 2.4GHz Wi-Fi 4 because of cost points and cost maturity of 2.4GHz only solutions for IoT and Camera applications.

From CommScope’s analysis of the home 802.11b and 802.11g only devices have almost fully flushed through to later and better Wi-Fi solutions.

There are several significant device types in the home that have different replacement cycles for Wi-Fi but typically group into the following categories

- Wi-Fi Access Point or Primary Gateway – typically Service Provider supplied
  - o Average time in the home – probably decreasing with technology acceleration but typically lasts for 5 years without change
  - o Typically, not changed by the SP when new technology introduced if the household has no issues or requires a new service that requires a GW or AP upgrade
  - o Wi-Fi upgrades may be tied to WAN access technology changes, such as DOCSIS 3.0 to DOCSIS 3.1 or GPON to XGS-PON
- Extenders – have a higher churn rate with consumers as they move through lower cost devices that underperform purchase improvement devices when necessary – probably a 3-4 year life cycle. Consumers have purchased retail repeaters themselves to remediate SP Wi-Fi issues. Some consumers have purchased their own meshing Wi-Fi solutions in recent times to bypass the primary service provider issued GW solution.
- Wi-Fi Client Devices
  - o Static Clients
    - Wi-Fi Smart TV – TV's have a 5-7 year life but can be moved from Room to Room
    - Wi-Fi STB – churn faster than Smart TV but typically have 3-4 year lifecycle and higher for quality SP issued solutions or Apple TV devices
    - Wi-Fi printers – 5 year cycles
    - Wi-Fi based PC/iMac – 5+ year cycles
    - Smart Assistants – 5+ year cycles
    - Camera typically indoor but often edge of range outside the house – Doorbell or Security cameras
  - o Mobile Clients
    - Smart Phones – depending on model and who supplies the phone but now tracking 2.5 to 2.88 years




**Figure 6 - Average time to replace Smartphone in the US**

- Tablets – until they break or die
- Laptops – 4-5 years. Corporate Laptops churn faster at 3 year cycle



- Apple Watch – churn time 3-4 years. Note even Apple Watch just uses 2.4GHz Wi-Fi.
- Low cost IoT Wi-Fi devices (typically last a long time until they die)
  - Smart Power Switches
  - Smart Sensors



**Figure 7 - Today's Wi-Fi architectures**

There are 3 typical Wi-Fi architectures found in most Cable Homes.

### 1. The most economical and the majority of households

Single all-in-one DOCSIS and Wi-Fi Gateway. Offers the lowest cost to the Cable Operator to provide both DOCSIS access and Wi-Fi services. Typically does a great job in homes less than 2,600sqft and offers also the highest MTBF because of single PSU. Highest Self install capability particularly in homes with existing DOCSIS modems as a replacement.

Downsides of the all-in-one Gateway is that it tends to be located at an outside wall location close to Cable ingress point to the home rather than more centrally in the home to accommodate better whole home Wi-Fi coverage.

### 2. The '2 box' solution where a separate eMTA and Wi-Fi AP are installed

Typically, both devices are connected by 3 feet of ethernet. Main driver for this architecture is to separate out the cadence of DOCSIS evolution and Wi-Fi evolution. The faster speed cadence of Wi-Fi technology and the consumer desire for the latest Wi-Fi solution has driven this solution. It does carry a higher initial CAPEX costs and probably over its lifetime, but the key metric is how often both the eMTA and AP are replaced vs one of them only. It also offers potential Operational improvements particularly for CSP's with multiple Access Networks like PON, DOCSIS, FWA where the same LAN based AP can be added to all access networks.

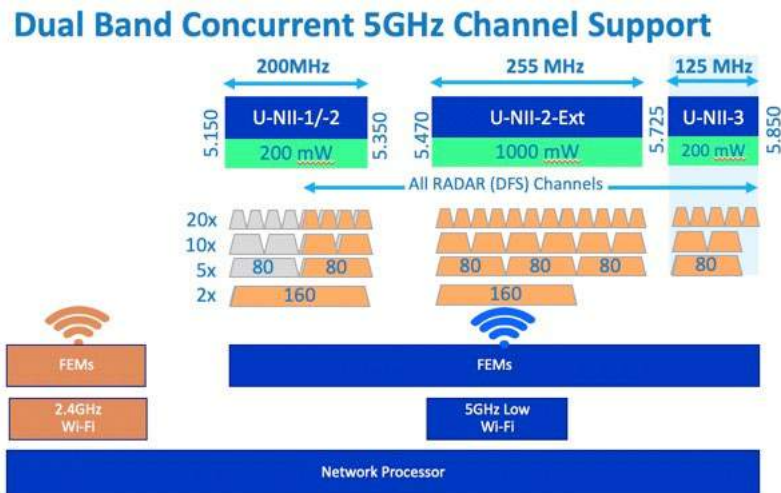
This architecture is also the Cable Operators nemesis when customers add their own Access point and NAT out the Cable Operator services.



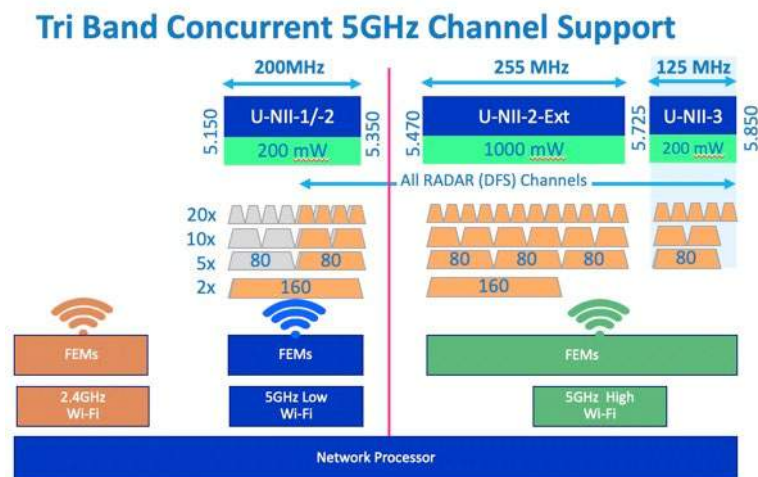
### 3. The Wi-Fi Mesh solution

We have now moved into the era of the Cable Operator deployed Mesh solution that extends Wi-Fi across the home, especially in larger homes. There are lots of discussions how best to extend coverage and the types of Wi-Fi extension

- Repeaters – trending away from these now to smart Repeaters
- Dual Band Smart Repeaters – employing some Wi-Fi Management RRM solution trying to maximize channel and band usage across the home
- Tri-Band extenders – with typically 5GHz backhaul solutions on a dedicated backhaul link that can also be used to attach clients under QoS mechanisms. Tri-band typically 2.4GHz+5GHz+5GHz – with U-NII1/2 and U-NII3 support one each radio.
- Other debates continue on how many streams and radios to use. Popular setups now are
  - 4x4 (2.4GHz) + 4x4 (5GHz) for the primary Wi-Fi AP when DBC
  - 4x4 (2.4GHz) + 4x4 (5GHz) + 4x4 (5GHz) for Wi-Fi AP when TBC
    - Similarly, for High End TBC extender
    - Lower end 4x4 (2.4GHz) + 4x4 (5GHz backhaul) + 2x2 (5GHz)
  - Low cost Smart Repeaters 2x2 (2.4GHz) + 2x2 (5GHz)



**Figure 8 - Dual Band Concurrent AP (2.4+5GHz)**




**Figure 9 - Tri Band Concurrent (2.4+5GHz Low+5GHz High)**

These have been introduced initially on Wi-Fi 5 architectures and within the constraints of 5GHz frequency availability in particular where 160MHz channel use has been almost non-existent because of (i) lack of client support during Wi-Fi 5 standard and (ii) no channels available with congestion and DFS requirements. Typically, most of the Wi-Fi transmissions in 5GHz are using 80MHz and 40MHz channels.

### 3. Service Providers and the introduction of Wi-Fi 6

Wi-Fi 6 introduces a significant number of new features for the operation of Wi-Fi, including a 4x average improvement per user in dense or congested environments (stadiums, etc.), a 4x improvement on network efficiency, and far better battery life for Wi-Fi 6 devices. Most importantly it offers a significant boost in performance over Wi-Fi 5 in terms of rate and range, promising about a 40% increase in data rate. The addition of 1024-QAM and 8 spatial streams as part of Wi-Fi 6 enables the highest speed of 9.6Gbps. From an AP technology perspective, it includes general signal handling improvements that not only benefit Wi-Fi 6 devices, but also previous generation devices, based on improvements with silicon process and front-end modules (FEMs).

Virtually all shipping Wi-Fi 6 APs are dual band concurrent, bringing these new benefits to both the 2.4 and 5GHz bands. Wi-Fi 5 did not address the 2.4GHz band, while Wi-Fi 6 has introduced improvements bringing beamforming and Multi User (MU) MIMO operation to 2.4GHz. Wi-fi 6 capability of using very narrow 2MHz frequency allocations to 2.4GHz clients enables a big increase in the range that is possible, as well as reducing power requirements for some IoT devices where they are not required to transmit across the entire minimum 20MHz channel as with previous versions of Wi-Fi. IoT devices can also benefit from Target Wakeup Time feature allowing them to spend more time “asleep” and save battery life.



**Figure 10 - The Primary Features of Wi-Fi 6**

The addition of OFDMA in Wi-Fi 6 brings a lot more control over how the allocated channel spectrum is used by devices and can be used for both Downlink and Uplink communication. Traditional OFDM clients are allocated channel access on the basis of time only, where a single client is assigned a period of time to use a channel. OFDMA is a multi-user version of OFDM, enabling control over both time and frequency enabling multiple users to transmit at the same time, using different/unique ranges of frequencies allocated to them. Wi-Fi 6 has quadrupled the number of sub-carriers within a 20MHz channel to 256, by reducing the sub-carrier width from 312.5KHz down to 78.125KHz. These sub-carriers are grouped together into smaller groups spanning contiguous spectrum, known as Resource Units (RUs). Wi-Fi 6 OFDMA operation assigns different RUs to different clients to use at the same time. OFDMA enables a reduction in the time a client must wait to transmit data, improving performance and latency in Wi-Fi 6.



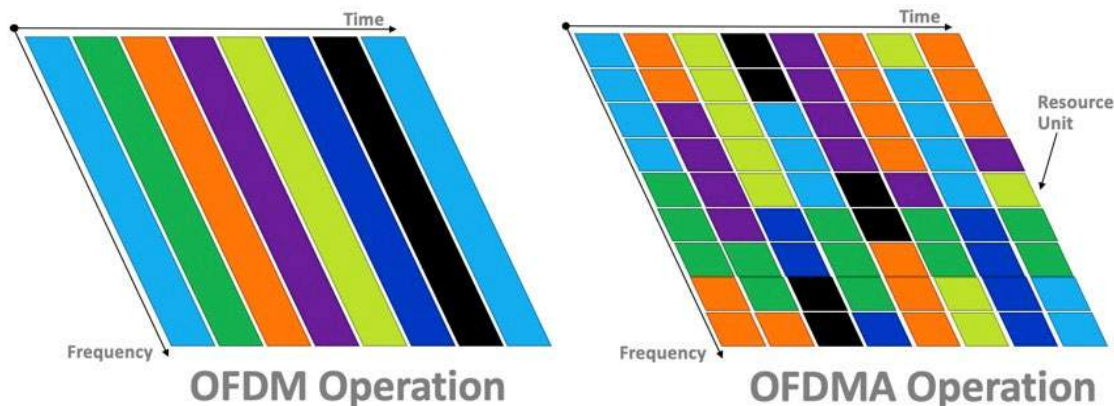
**Figure 11 - 20MHz RU allocation mapping**



**Figure 12 - 40MHz RU allocation mapping**

The ability to have parallel transmissions means that applications, such as audio or video streaming, that send chunks of data continuously can be served simultaneously. A 20MHz channel can support up to 9 concurrent clients, while a 160MHz channel can transmit to 74 clients in parallel. mixtures of RU allocations are also possible, providing more spectrum to some clients than others, depending on airtime demand and policies.


The AP is responsible for downlink OFDMA transmissions, coordinating the client to RU allocations. For uplink OFDMA transmissions, the AP coordinates the different associated clients and issues a trigger frame to get them to transmit at the same time. All clients using OFDMA must tune their power, frequency and timing to be able to join in the parallel transmission



As can be seen from the above diagram, the ability to send or receive traffic from different Wi-Fi users in the same time period has a significant improvement in link sharing and latency reduction. The “dark blue” user above gets services in the 3<sup>rd</sup> time period using OFDMA compared to the 7<sup>th</sup> time period with OFDM. There is a bit more signaling required for OFDMA in terms of link overhead, but the benefits, especially when user devices are sending short bursts of traffic (typical for interactive applications, or even voice/video conferencing), outweigh these significantly. Having parallel transmissions, as opposed to consecutive transmissions enables much lower latency which brings benefit to all services. A 75% reduction in latency is possible with Wi-Fi 6, due to OFDMA operating in both the DS and the US. Coexistence with Wi-Fi 5, where the traditional airtime access procedure (EDCA access) is present can be challenging, but through effective management of the EDCA minimum contention window ( $W_0$ ) and maximum contention stage ( $m$ ), controlled access between Wi-Fi 6 and Wi-Fi 5 devices can be achieved.

Wi-Fi 6 also supports up to 8 spatial streams, (Wi-Fi 5 offered it already but was rarely implemented) along with Multi User (MU)-MIMO. Most Wi-Fi clients include 1 or 2 spatial streams, so with an 8x8 AP, it will be possible to service 4 such clients at the same time using MU-MIMO, delivering gigabit speeds to each device. DL MU-MIMO is a mandatory feature of Wi-Fi 6 certification. Depending on the location of an AP relative to clients within a home, the use of MU-MIMO can really improve networking performance. Note that MU-MIMO is very effective at short to medium ranges, while OFDMA is effective at all ranges, short, medium and far. Also note that MU-MIMO is also available for 2.4GHz band, something that was not part of the Wi-Fi 5 standard. And finally note that it is possible to operate OFDMA and DL MU-MIMO together, identifying different users to different spatial streams while also using OFDMA to offer partial bandwidth to users on the same spatial streams.

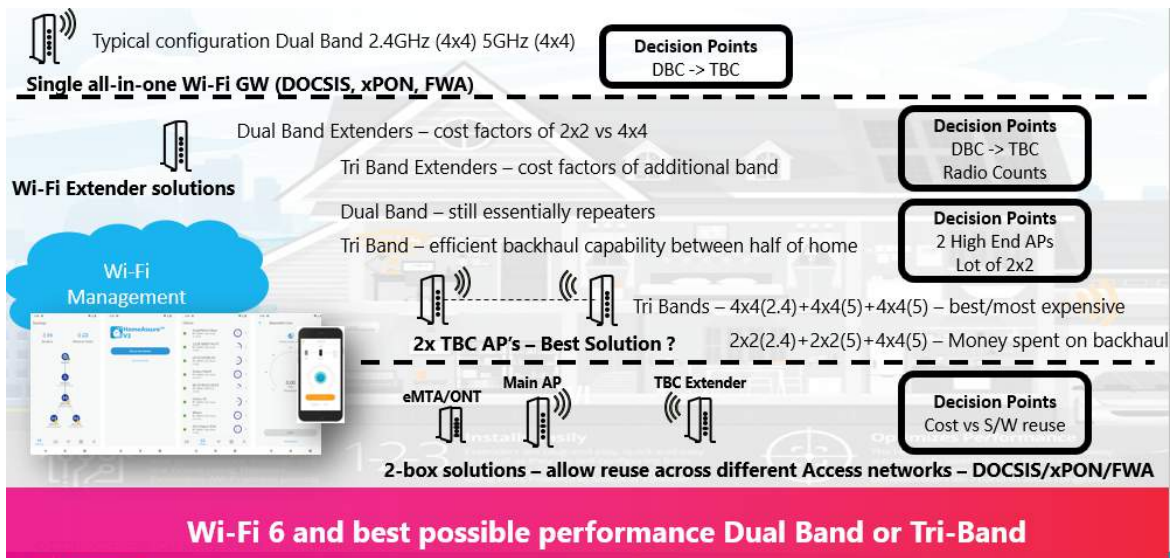




**Figure 13 - DL MU-MIMO + OFDMA**


The above figure shows how DL MU-MIMO and DL OFDMA can be put to good use, enabling 7 different users to receive data within the same transmission period. One user (in light green) received 3 separate RUs (106 sub-carriers each) from 3 different spatial streams, 5 other users got individual 106 sub-carrier RUs, while 1 user got 4 separate 26 sub-carrier RUs, one from each spatial stream. With increasing bandwidth to 160MHz, operating with a 4x4 setup, OFDMA+DL MU-MIMO will really improve network performance.

#### Decision points on Wi-Fi 6 introduction



**Figure 14 The decision points for Wi-Fi 6 introduction**

From a service provider perspective, the introduction of a Wi-Fi 6 AP brings performance benefits even for previous Wi-Fi generations due to the hardware design improvements, including range and rate.

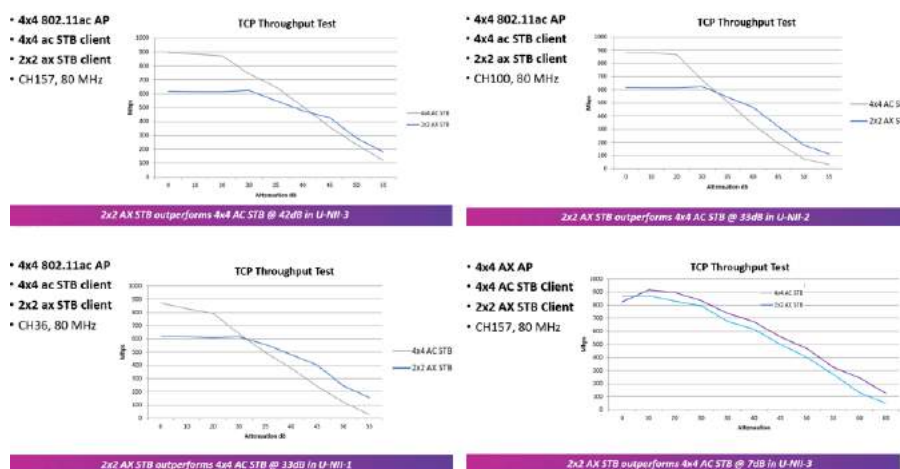


**Figure 15 - Wi-Fi 5 vs Wi-Fi 6 2x2 Client Performance Capability (Mbps)**


The real benefits are only realized when new Wi-Fi 6 non-AP STAs (devices) are added, such as service provider supplied STBs/SMDs and mesh extenders, or consumer owned mobile devices (phones, tablets) or media streamers (smart TVs, Amazon fire, etc.). The mix of Wi-Fi 6 APs matched with Wi-Fi 6 devices enables all the new technology features in Wi-Fi 6. See below the performance of

- AC only AP with an AX STB client – showing improved performance at range of the AX client with the AC AP
- AX AP against AC and AX STB – showing improved performance for both

**CommScope recommends that future STB should use 2x2 instead of 4x4 to get the performance needed at lowest cost points and lowest thermals.**



**Figure 16 - STB client performance showing 2x2 Wi-Fi 6 outperforming AC 4x4**



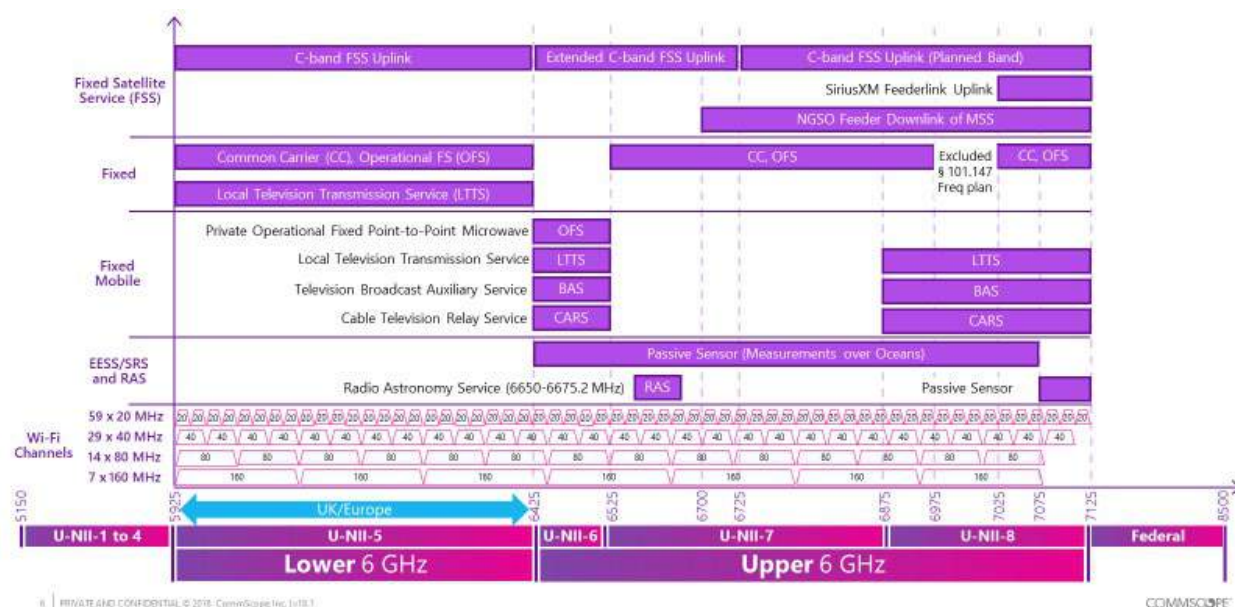
**Figure 17 - STB client performance showing 2x2 Wi-Fi 6 outperforming AC 4x4**

STB devices with Wi-Fi 6 will work extremely well with Wi-Fi 6 APs, bringing better overall experience to subscribers from greater range as well as more reliable service. Wi-Fi Extenders used to add coverage throughout a subscriber home will also benefit, especially due to the lower latency provided by Wi-Fi 6, which becomes very important when data has to bounce from the AP to the Extender, and then from an extender to a Wi-Fi device.

The ability to use OFDMA for serving AP connected devices and extenders in parallel cuts latency significantly when serving devices connected to extenders, as any queued traffic can be sent at the same time. Multiple hops back and forth through extenders all increase latency, which then has a knock on effect on the throughput devices can achieve – lower the latency and throughput increases, enabling adaptive streaming protocols like DASH/HLS/etc. used by STBs to operate at higher display resolutions, etc., giving a better quality of experience to the end subscriber.

#### 4. Setting the stage for Wi-Fi 7 entrance – Wi-Fi 6E

The available spectrum used in Wi-Fi 6 is the same as Wi-Fi 5, and is limited to long standing 2.4 and 5GHz regulatory bands defined by the FCC, which covers about 400MHz of spectrum. Recent changes by the FCC, after many years of discussion between industry and regulatory stakeholders has led to the opening up of the 6GHz band offering up to 1.2GHz of additional spectrum for unlicensed usage, with Wi-Fi being one of the key benefactors of this spectrum. This effectively quadruples the spectrum available for Wi-Fi. To help differentiate the use of the 6GHz spectrum Wi-Fi Alliance have coined the term “Wi-Fi 6E” to identify Wi-Fi 6 devices capable of operating in this “Extended” 6GHz spectrum.



**Figure 18 - US 6GHz spectrum and channel allocation to Wi-Fi**

This extra spectrum will enable Wi-Fi innovation for decades to come, and some reports indicate it will add over \$183B to the US economy by 2025, with Wi-Fi equipment sales being \$1.54B of that figure. A key element of the use of 6GHz is the displacement of the need to adopt 5G cellular connections for high speed wireless access, leading to significant cost savings of over \$54B by enterprise customers (note that this may be impacted by consequences of COVID-19). With the extra spectrum capacity, Wi-Fi has room to grow, and most users of Wi-Fi would likely stick with Wi-Fi than need to switch to a completely new 5G architecture, be it licensed or unlicensed.

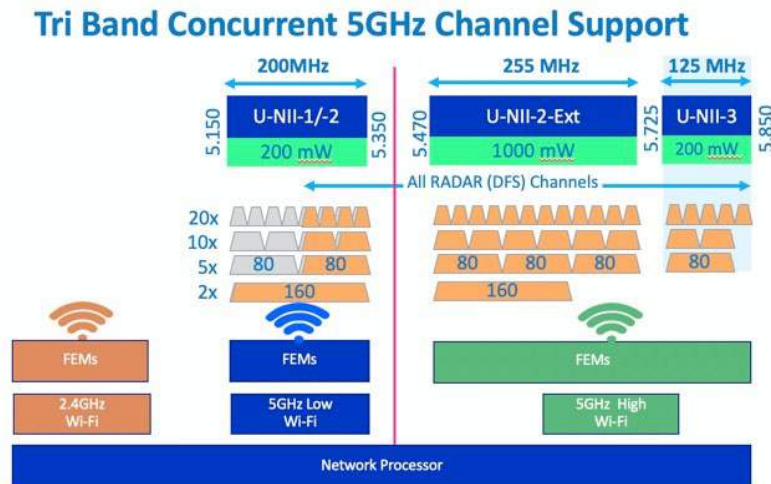
One of the biggest drawbacks with Wi-Fi has been the lack of available spectrum to enable the HW capability of Wi-Fi 5 and Wi-Fi 6 APs that offered 80 and 160MHz channels. The existing spectrum enabled a maximum of two 160MHz channels, which typically required a TBC (Tri Band Concurrent) AP enabled with two 5GHz radio lineups. However constraints in the spectrum related to DFS/Radar access meant that it was often not possible to get that capacity, and APs were limited to using the available 80MHz wide channels (additional constraints relating to overlapping Wi-Fi APs in dense spaces also meant that 160MHz channels were difficult to operate due to interference).

|                    | 2.4 GHz | 5 GHz   | 6 GHz     |
|--------------------|---------|---------|-----------|
| Available Spectrum | 85 MHz  | 480 MHz | 1,200 MHz |
| 20 MHz             | 3       | 25      | 59        |
| 40 MHz             | 2       | 12      | 29        |
| 80 MHz             | 0       | 6       | 14        |
| 160 MHz            | 0       | 2       | 7         |
| 320 MHz            | 0       | 1       | 3         |

**Figure 19 - Channel Availability over Spectrum**



The 6GHz spectrum that the FCC has released has the capacity to support 7 additional 160MHz channels in addition to the existing roughly 400MHz of capacity in 2.4 and 5GHz. A single 160MHz Wi-Fi 6 channel can support up to 9.6Gbps (when using 8 spatial streams), meaning that the 6GHz band can offer over 67 Gbps (7 \* 9.6 Gbps) of additional bandwidth.




**Figure 20 - TBC 2.4+5GHz AP**

The 1.2GHz of spectrum in 6GHz was already categorized by the FCC into 4 separate bands, that were aligned to the various requirements of existing incumbent users. In order to operate in this new spectrum, APs either have to lower their operating power or work with a spectrum supervisory service

Wi-Fi 6e builds on top of Wi-Fi 6 to enable operation in the 6GHz spectrum, identifying the necessary controls around how the spectrum is channelized and various performance limits. Channel access details are also included to ensure a common set of rules around how 6GHz APs and devices can use the new spectrum. Wi-Fi 6e will be limited to the current performance outlined in the Wi-Fi 6 standard (which is 9.6Gbps).

The following diagram outlines the UDP rate over range/attenuation for a 4x4 AP to a 4x4 Client (typically a backhaul application on a 6E mesh solution). This is done at the current 5dBm/MHz PSD approved by the recent FCC Low Power Indoor specifications. **Note – all the Throughput/Attenuation plots in this paper assume perfect signal, with no noise and no co channel interference. They also show just Tx potential only and don't take into account factors like the Wi-Fi beacons which transmit in 20MHz channels and at 18dBm EIRP. So, the actual performance will be somewhat poorer than illustrated. Unless otherwise stated as a specific test for different client configurations assume all plots are just theoretical best case from AP for illustrative purposes.**



**Figure 21 - 4x4 to 4x4 Wi-Fi 6E performance at 5dBm/PSD - LPI specification**

In addition to the Wi-Fi Alliances efforts on enabling access to the 6E spectrum using Wi-Fi, one of the main questions affecting equipment vendors is how to create cost effective hardware that addresses the existing 2.4 and 5GHz landscape while also offering the benefits of 6GHz

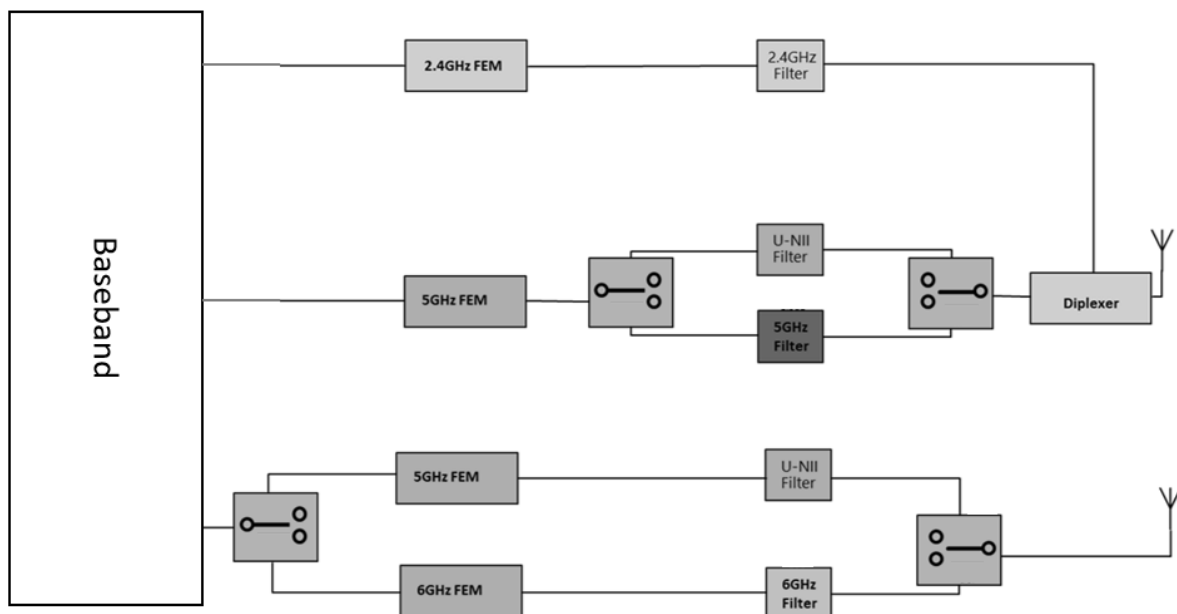
Current TBC designs that operate two 5GHz radio lineups (5GHz Low (UNII-1 and UNII-2) and 5GHz High (UNII-2 Extended and UNII-3)) need additional filtering and power in order to provide the best performance. Operating two transceivers with just 160MHz of separation has required considerable isolation filtering to ensure that as one transceiver transmits at a high power that the other transceiver can receive signals from connected devices. The filtering adds cost to the standard design. Another constraint is that some Wi-Fi devices limited themselves to only operate in non-DFS channels, which is an 80MHz wide part of the 5GHz Low spectrum, meaning that this part of the spectrum must be supported.

Including a separate dedicated 6GHz only subsystem into an AP is possible from an engineering standpoint but adds additional cost to a product that may see limited initial 6GHz use while waiting for new devices to appear. Such a solution may also increase the industrial design/enclosure size beyond what would be acceptable for a home Wi-Fi AP format factor, as it is likely there would be 15 radios and associated antenna (3+4+4+4) for the 4 different bands to be supported. Such a unit is unlikely at this point in time. The simplest and likeliest deployment of 6GHz is to offer a tri-band that supports 2GHz+5GHz+6GHz. The tradeoff as discussed is the additional complexity of supporting U-NII1/2/3 in the 5GHz band and the initial relative little use of the 6GHz radio with the sparsity of 6E clients in the home.

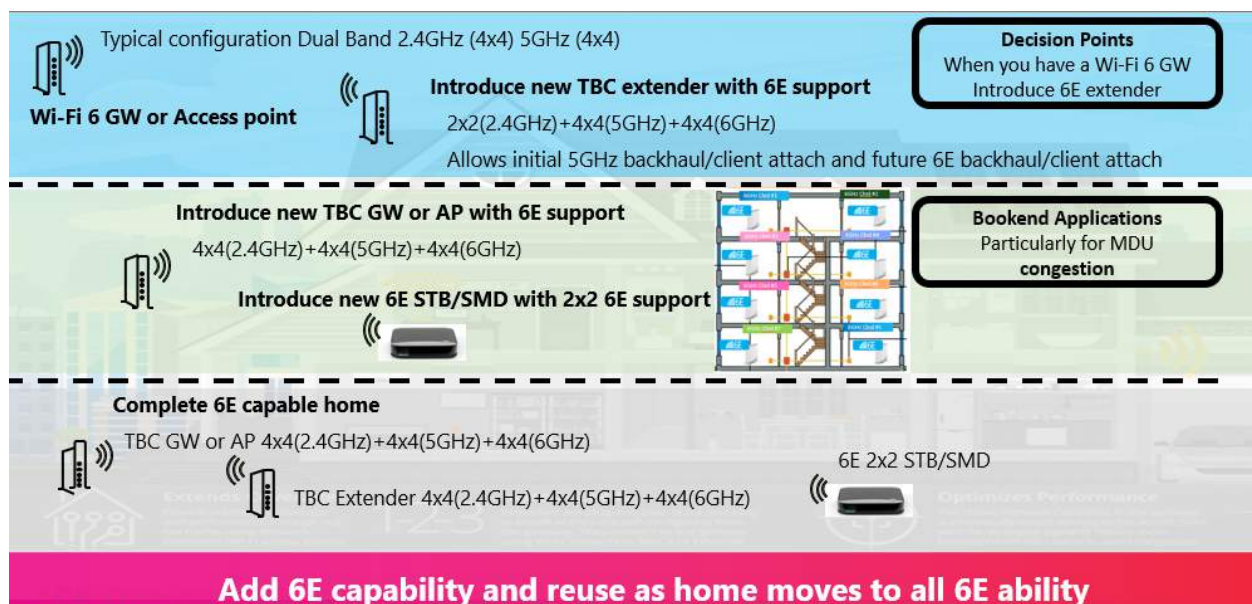
Options are being considered as to what makes the best combination of low cost and flexibility for radio operation across the 5/6GHz space. Note that flexibility in terms of switching between filters and radio lineups unfortunately has a knock on effect on the total output power, as signal losses creep in as a result of passing signals through switch banks that enable such flexibility. Switching solutions overall also significantly increase the cost of the AP so for the moment its unlikely that Cable Operator devices will add 5/6 switching on the radio. Additionally, the decision as to when to switch a radio to use either 5GHz

or 6GHz is also complex. Manually switched based on a consumer or Cable Operator policy or some algorithm based on some intermittent sampling of the 6GHz capabilities of the home 6GHz client arrivals. The following diagram illustrates the complexity of trying to support a Tri-band that can

- Switch between all U-NII1/2/3 bands on 2 chains
- Switch between 5GHz and 6GHz
- Support the Filtering required for 5/6GHz and within the 6GHz channel plan



**Figure 22 - Switchable 5/6 designs will bring in cost and complexity levels that may not be acceptable to Operator cost points**



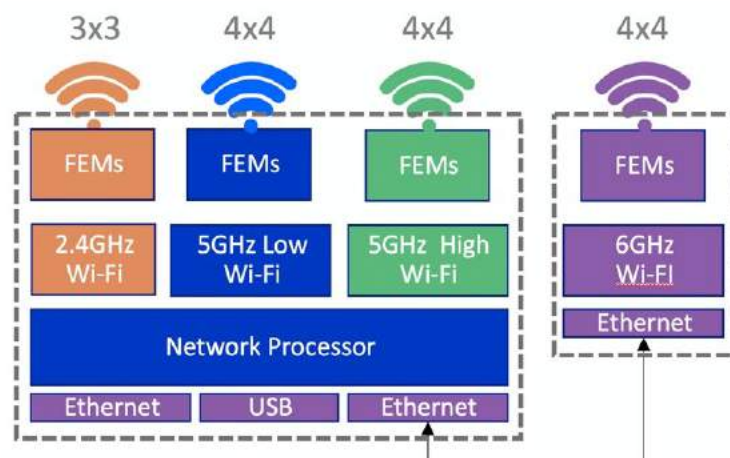
The 5GHz High spectrum is closest to the new 6GHz spectrum, so there is some interest in having a TBC AP design that continues to offer 2.4 and 5GHz Low and mixed 5GHz High and 6GHz support. This appears like a good tradeoff, but a key drawback with this approach is that the 5GHz High spectrum that is subject to DFS is also a part of the spectrum that enables greater radio output power (1W compared to 200mW in the other 5GHz Low spectrum). Using a mixed 5GHz/6GHz option means that once a single 6GHz device associates with the AP, the 5GHz High spectrum is no longer available as the AP has to switchover to 6GHz.

A lot of recent AP designs have incorporated Wi-Fi 6, and are either being deployed or going through the final stages of acceptance testing in service provider labs. Obviously, service providers want to keep in step with technology choices, but struggle with the economics and operational headaches of deploying a new Wi-Fi 6E AP so soon after launching their Wi-Fi 6 AP. To overcome this issue and enabling the adoption of 6GHz, one approach is to use 6GHz adapter devices. Ethernet (either 1Gbps or 2.5GBps) is used to connect to these new 6GHz adapters. Such an option is not unique, as pre-standard Wi-Fi 5 video bridge adapters were deployed by many operators in the past to provide whole home HD video streaming over Wi-Fi to ethernet connected IP STBs.



**Figure 23 - ARRIS VAP2400 Wi-Fi Video Bridge**

This approach relies on one adapter being connected to the main AP, configured to create a new 6GHz Wi-Fi network that other devices connect. These 6GHz adapters enable the adoption of 6GHz within the home, either to act as a 6GHz Wi-Fi backhaul to existing extenders in the home, or to enable subscriber devices take advantage of the high speed, low latency wireless connectivity over 6GHz back to the main AP. No change is required on the existing home devices, apart from using Ethernet to connect to the adapters. USB is also a possible interface (offering both a data/control path as well as bus powering).



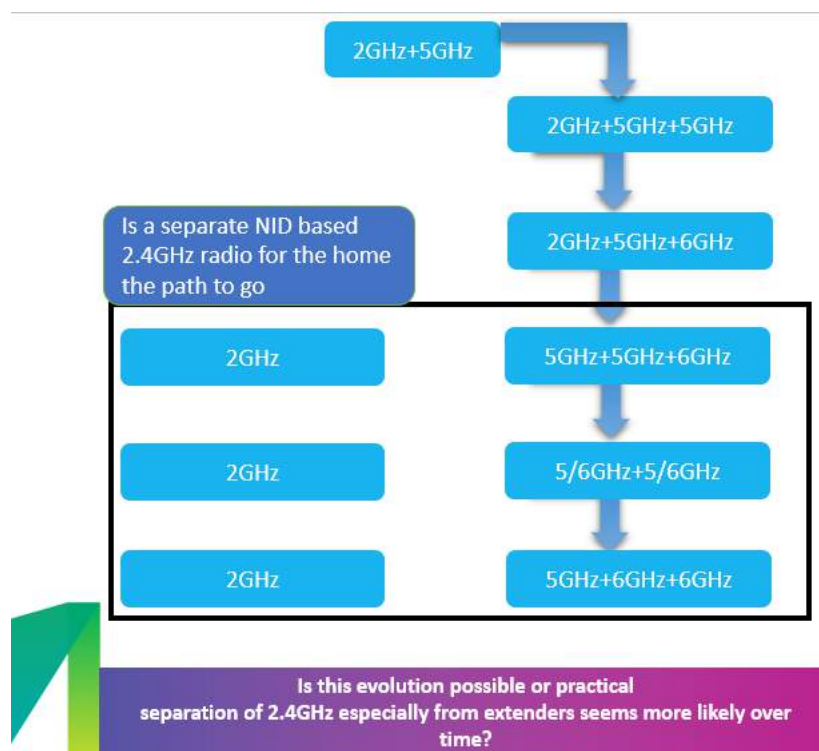
**Figure 24 - External Wi-Fi 6E dongle for existing GW**

This also brings up the question of the effective relevance of the 2.4GHz radio. With its primary purpose today as

- Supporting low cost devices like cameras and IoT devices
- Providing the comfort of connectivity at range particularly in single AP homes without meshing systems extending 5GHz particularly.

But given the sheer capacity of the spectrum shift now to 6GHz which has only 80MHz – 450-600Mbps vs 66Gbps capacity in the 6GHz band with 8 spatial streams and the Wi-Fi 7 standard even allowing for 16 spatial streams – what role does 2.4GHz have going forward in the home Wi-Fi ecosystem. Should we now look to

- Have a single point in the home for 2.4GHz AP for legacy devices and to support 2MHz channel IoT capabilities built in the Wi-Fi 6 standard for up to 4 times more range on narrow carriers.
- Focus a 5GHz and 6GHz solution on the high capacity single AP across the home or the multiple mesh APs – for maximized use of higher capacity bands and take every opportunity to reduce any effect on 5/6 designs that 2.4GHz drags in and see can we
  - o Lower gate counts in silicon
  - o Not having to trade off antenna placement for best 2.4GHz and 5GHz on the same antenna
  - o lower costs and reduce thermals



**Figure 25 - Can 2.4GHz be separated from the Wi-Fi mesh 5/6/7 devices in some future architecture**

It may take until Wi-Fi 7 is introduced to relook at the grouping of 2.4GHz the 5GHz U-NII-1/2/3 bands and the U-NII5/6/7/8 bands for optimal architectures in the home based on the evolution of FEM and



Filter technology to try and ensure the best possible designs to address the widest frequency range in the smallest form factor package.

This idea may not have a lot going for it, but in some ways the other 2.4GHz wireless systems, such as Bluetooth and Zigbee, putting all of these together into its own AP package may make sense for optimal coverage given that 2MHz use of 2.4GHz, Zigbee and BLE would have similar ranges from the same location of this 2.4GHz AP as illustrated below

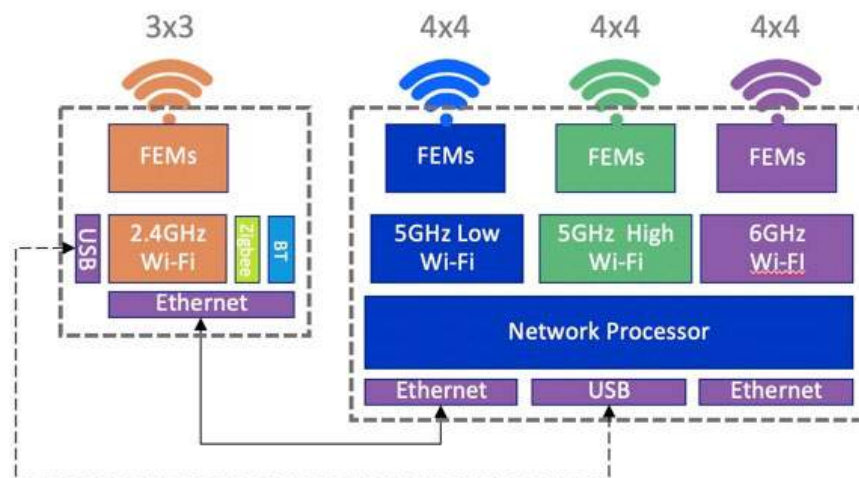


Figure 26 - Add 6GHz to AP, but extraction all 2.4GHz Radios


## 5. Low Power Indoor Mode to AFC controller Standard Power

Some of the FCC rules about 6GHz operation are specifically about power levels that APs can use when operating in the band. Serious consideration was given to the coexistence of future unlicensed operation of Wi-Fi devices and the incumbent 6GHz users (mostly Microwave links, Fixed Satellite Services (FSS), Cable TV relay and Broadcast services). This has led to various compromises about acceptable power limits across the different UNII bands.

A Low Power Indoor (LPI) only power level applies across the entire 6GHz band (from the bottom limit of UNII-5 to the upper edge of the UNII-8 band). This corresponds to a Power Spectral Density of 5dBm/MHz EIRP. For a 20MHz channel this is 18 dBm, while a 160MHz channel is allowed 27 dBm. Note that LPI APs can only operate indoors, and the FCC rules have very explicit limitations for such operation, as follows.

- Low-power APs can't be weather resistant, preventing them from being installed outside permanently;
- Low-power APs can only have integrated antennas and it is prohibited to provide the capability of connecting other antennas to the devices;
- Low-power APs can't be battery operated;
- Low-power APs must be clearly labelled with a notice that states that "FCC regulations restrict operation of this device to indoor use only."

From a 6GHz regulatory perspective in the US the following is in play to try and number the potential channels in the 1.2GHz of 6GHz as well as define which channels and bands can be used for Standard Power under AFC control.



**Figure 27 - Current potential channel numbering scheme for U-NII5 to U-NII8**

Standard Power (SP) is a higher output power that APs can also use but under very strict conditions. Such power levels are only applicable to UNII-5 and UNII-7 bands (comprising 850MHz of the total of 1.2GHz of spectrum). SP operation is subject to a spectrum supervisory service called Automated Frequency Coordination (AFC). The AFC system relies on the AP reporting its geographical location to the AFC to identify any incumbent systems in the area, and what frequencies those systems may be using. The AFC will make a decision if Wi-Fi 6E operation is allowed by the AP, and if so, will identify the exact list of frequencies and power limits that apply. The AP must repeat this exchange with the AFC on a regular basis, as the data that the AFC relies upon may receive dynamic updates relating to existing or new 6GHz users that could impact the APs use of 6GHz. This approach will enable coexistence with the incumbent users of this spectrum. 320MHz channels are now part of the IEEE802.11be initiative and there is no final answer yet. The only 2x160MHz channels that fit in U-NII band is U-NII-5. There is also some opportunity to put 320MHz channel in U-NII-7 but as its mis-aligned in the current 160MHz allocation discussions are ongoing how this could be managed

Once an AP is enabled and compliant with the AFC system it is allowed operate outdoors as well as at the higher power level of 36 dBm EIRP for all channel bandwidths. This is a key enabler for outdoor APs, as well as potentially enabling higher indoor powers. As an industry we are still working through the architectures and process to be able to give location to an Indoor AP particularly and there are various mechanisms being considered that are trying to exclude the use of any additional GPS cost burdens on the AP. The frequency of the AP having to check with the AFC is also another part of the equation to try and create a cost efficient but friendly solution to incumbents. More on this area as the decisions unfold.

## 6. Wi-Fi 7 and the 320MHz channel what does it mean for the home

The current highest speed possible with 802.11ax on a 160MHz wide channel is 1.2Gbps per spatial stream. Using the typical allocation of 4 spatial streams ~4.8Gbps can be realized and with the maximum of 8 spatial streams in MIMO, this can enable a maximum data rate of just over 9.6Gbps, and this requires both the AP and client device to support 8x8 operation. The IEEE802.11be is now proposing to add the following features to the Wi-Fi 6E specification

- 320MHz and more efficient utilization of non-contiguous spectrum
- Multi-band/Multi-channel aggregation and operation
  - Multi-Link Operation - MLO. Arguably one of the most important features from a pure BW or latency perspective. MLO also ties together the important of 2.4GHz, 5GHz and 6 GHz frequencies and bands as you can utilize all or any combination of the bands
- 16 spatial streams and MIMO protocols enhancements
- Multi-AP coordination (e.g. Coordinated and Joint Transmission)
- Enhanced Link adaptation and transmission protocol (e.g. HARQ)
- Adaptation to regulatory rules specific to 6GHz spectrum
- Refinement of 802.11ax features

If there are two clear 160MHz channels in 5GHz, and an AP supports 8x8 for both 5GHz Low and 5GHz High, this can allow for up to 19.2Gbps of capacity at just 8 spatial streams. In the next version of the Wi-Fi standard, 802.11be, the primary feature that is being added is the ability to use a 320MHz channel and to increase the spatial streams to 16 for the current 8 in Wi-Fi 6. This fundamentally then creates a unique platform to feed the new low latency and future high bit rate services. In the US with the entire 6GHz spectrum being 33-66Gbps of capacity it creates a fundamental platform for low latency and high capacity services. This paper is not going to outline these services but for example LightField TV at 8K levels comprising 43 or more planes of video at 8K rates can reach well above 2Gbps levels. If you look at products from Looking Glass ([www.lookingglassfactory.com](http://www.lookingglassfactory.com)) you will see 'holographic' LightField TV solutions giving 'no glasses required' immersive viewing experience and showing real depth and volume. Solutions like Wi-Fi 7 and Wi-Fi 6E will enable these solutions to be un-tethered from Ethernet interfaces at 2.5Gbps and even 10Gbps.






**Figure 28 - Immersive display on the LightField Video Looking Glass display at 8K levels**

If we look at the current FCC specification for 320MHz channel usage in the 6GHz band, the power level has been set to 30dBm for the 320MHz channel. With 30dBm at 5dBm/MHz the following performance levels can be expected. You can see below the performance of both 8x8 and 4x4 in the UNII-5 band with 30dBm. Performance levels are substantially improved with Airtime throughput over range and offer the potential of really good whole home coverage from a single GW/AP for all but the largest of homes in the US. Can we realize the following architecture for Wi-Fi 7 – with single high performance AP driving Gbps across the home. Up to 20Gbps could be realized with an 8x8 device with 8 spatial streams.




**Figure 29 - Basic single 30dBm - 8x8 or 4x4 Cross Home Wi-Fi 7 Architecture**

This single 1W 8x8 device with only 8 spatial streams providing cross home Gbps speeds driving opportunity again for single AP and no ethernet wire solutions for homes in the 2600sqft to 5,000sqft range. The rates range close to theoretic 20Gbps at the AP to maintaining a Gbps to 54dB of attenuation.




**Figure 30 - 8x8 AP/GW at 30dBm in both 320MHz and 160MHz channels**

Even for a 4x4 the potential for high performance across the home is high with a Gbps of throughput potential even at 55dB of loss and theoretically ~9Gbps at the Access point.




**Figure 31 - 4x4 performance at 320MHz and 160MHz at 30dBm**

As AFC standard power currently allows for higher power solutions outside – if we were to have approvals for 33dBm/2W Radiated power in a solution its performance over attenuation would improve as defined in following graphs.



**Figure 32 - 8x8 performance of 320MHz and 160MHz at 2W radiated Power**

At 33dBm GBps speeds would be maintained to almost 54dB attenuation.




**Figure 33 - 4x4 performance of 320Mhz and 160MHz at 2W radiated Power**

Even 4x4 and 4 spatial streams would support Gbps speed to 53dBm attenuation.

However, there could be even more performance for Wi-Fi if we consider the room vs the home. With 6GHz spectrum we have the potential to create an intra room AP backbone using one 320MHz and an inter room 160MHz channel all using different spectrum and potentially to optimal single use levels particularly for larger homes.

If we assume, we want to bias more to 6GHz capable devices in some future scenario

- One primary home AP providing 320MHz cross home 8x8:8 (20Gbps) or 4x4 (10Gbps) at 30dBm.
  - o Wi-Fi 6E/7 capable clients can attach to this AP, but its primary role is providing intra room Backhaul
  - o The device could also offer 5GHz and 2.4GHz support as well
    - 4x4(2.4GHz) + 4x4 (5GHz) + 8x8 (6GHz) or 4x4(2.4GHz) + 4x4 (5GHz) + 4x4 (6GHz)
- In Room (potentially 4 other 160MHz channels that can be used – 4 rooms) that have a dual or triband
  - o This device could be a triband 4x4 (5GHz inroom) + 4x4 (6GHz in room) + 4x4 (6GHz backhaul)
    - Could potentially even use 25mW VLPI mode for in Room LAN to create smaller form factor and lower thermals. Device could potentially adapt power to the room conditions and increase output power with appropriate PA designs.




**Figure 34 - Performance of VLPI in room potential low power solution**

- o In room clients would also reduce uplink power to connect with these in room AP
- o The device could offer additional radios for 2.4GHz and 5GHz as well as 6GHz as backhaul



**Figure 35 - Wi-Fi 7 backhaul - Wi-Fi 6E in room - Ultimate Home Wi-Fi ?**

This is simply illustrated below with 3 in-room AP's using different in Room 160MHz channel to the 320MHz backbone channel across the home.




**Figure 36 - use of Wi-Fi backbone and in room 4.7GBps networks**

What would such an architecture be used for and what does 4.7GBps of in room clean Wi-Fi drive for new platform for applications and services. We are still trying to figure that out – but having this new bandwidth will accelerate more immersive video services and potentially drive applications like lower



latency VR bandwidth to Wi-Fi 7 capable HMDs using lower power Wi-Fi to make the connection in room, increase battery life of HMD and reduce Heat in the device.




**Figure 37 - Wi-Fi 7 and Wi-Fi 6E - driving high capacity low latency services**

Additional services with lower decoding or encoding time for latency sensitive video could emerge and certainly the bandwidths supported in 6GHz allow for wireless HDMI services for the first time on Wi-Fi above 4K or even 8K levels with the full 16 spatial stream capabilities that 802.11be is driving.

The above architecture is a North start where we can take this new home Wi-Fi capacity. Challenges exist to continue to

- Make the solutions cost effective for the masses vs the minority of home budgets
- Keep the size and thermals in check – as the ergonomics of Wi-Fi AP's still has a big factor on any architecture.
- Keep latency over Wi-Fi at a minimum to ensure latency sensitive services are covered

**The above requirements could see 320MHz solutions offered in a 2x2 solution play for 6GHz – given the math of getting to >4Gbps of airtime capacity on a 320Mhz channel. This would be in contrast to the higher more future proof capacity architectures at 8x8 and 4x4 illustrated above. So, there could be architectures that align more with keeping 4x4 for the 5GHz bands and legacy and using 2x2 for 320Mhz for 6GHz capable clients using Wi-Fi 7. Its not even clear if any initial Silicon directions would focus on a cost optimized 2x2 320Mhz solution for 6GHz.**



**Figure 38 - 2x2 320MHz Wi-Fi 7 vs 4x4 160MHz Wi-Fi 6E**

so, it will be interesting to see how single Wi-Fi 7 Wi-Fi Gateways and Access Points emerge to push into the highest performance levels vs keeping the cost economies for Cable Operators in particular. Who knows even Wi-Fi repeaters may make a comeback. Repeaters have always been problematic particularly for 2.4GHz low capacity or high congestion 5GHz networks as they receive and transmit the same packets on the same frequency typically halving the available airtime (typically more). With 6GHz channels however you have lots of capacity to half each time you repeat from 20Gbps in Wi-Fi 7 320MHz channels to ~4.8Gbps in 160MHz 6GHz channels.

A single Wi-Fi 7 Gateway with 320MHz capability and Tri-band support will probably be the primary workhorse for the Cable Operator from 2024/25 onwards. However, as guaranteed Gbps applications in room emerge there may be the first in room high capacity APs that potentially focus specifically on driving applications only for 6GHz capable clients.

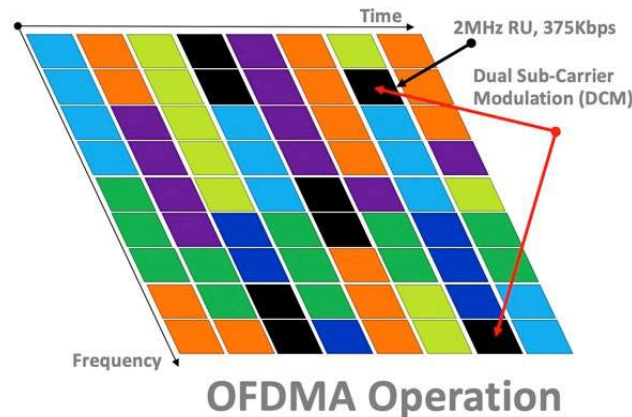
## 7. What do we do with 2.4GHz and 5GHz spectrum in the Wi-Fi 7 era

We have discussed a lot about the high capacity applications that can run in the 6GHz spectrum and why not with 66Gbps+ of capability there (increased spatial streams and modulation). Like any investment the focus should go where the growth will be so where does this leave the 2.4GHz spectrum and in particular the different performance/range ratio to 6GHz based devices. When 802.11ax was conceived one of the features added to it was to potentially complete the journey of 2.4GHz spectrum as the range work horse and the most likely to make a connection service, was a 2MHz channel that could be used for Narrow Band IoT services. This channel could carry the low bit rate IoT device services at ranges up to 4 times that today of 2.4GHz channels. This could then create a wider separation to the 6GHz access point role and would mean that a single 2.4GHz AP for NB-IOT functions in even the largest homes is feasible.

Wi-Fi 6 has brought fresh ideas to the 2.4GHz space, including the ability to deal with 2MHz wide channels for power sensitive IoT devices. Combining 2MHz, OFDMA, and Target Wakeup Time (TWT) together means that IoT devices are now capable of offering a significant competition to Zigbee, Z-Wave and BT devices. The 2MHz wide channel offers an 8 dB improvement on noise power, allowing the signal power to be 8 db lower enabling a greater coverage area for IoT devices. The TWT feature enables devices to remain in deep sleep mode for much longer than ever before, avoiding the constant need to listen to beacons transmitted by the AP, and allowing the conservation of the all important device battery life.



The flexible guard interval ranging from 0.8us (indoor) to 3.2us (outdoor) also makes the transmitted signals more robust, improving decoding ability at the receiver side. In addition to competing against existing LWPAN networks, Wi-Fi 6 also has the potential to avoid transmitting in the frequencies already used by these networks by not using OFDMA RUs that overlap the same frequencies. Another Wi-Fi 6 feature that helps IOT is the ability within OFDMA to repeat a device transmission in two different RUs, where a transmitted RU is repeated within the same time but on a different RU boundary.



**Figure 39 - OFDMA with 2MHz DCM**

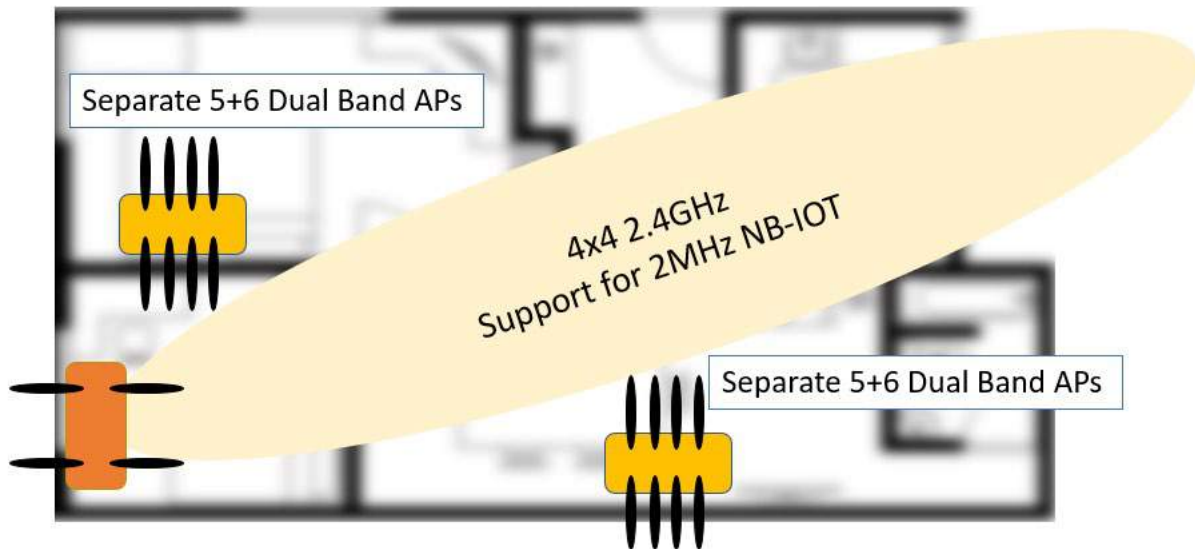
This Dual SubCarrier Modulation (DCM) feature helps the robustness of the IoT transmissions, increasing the chance of a successful data exchange. The use of OFDMA and 2MHz RUs (enabling 375Kbps) reduces the amount of bandwidth wasted by previous versions of Wi-Fi that had a minimum speed of about 6.5Mbps, and allows other active users to operate at the same time



**Figure 40 - Improved Wi-Fi 6 IoT co-existence with ZB/BT**

The 5GHz band is certainly getting close to capacity with existing Wi-Fi 4 and 5 devices, most of which are likely to remain in use for an extremely long time. While cellphones get replaced at some regularity, most other Wi-Fi devices, as long as they perform satisfactorily, are unlikely to be switched out. The gradual introduction of Wi-Fi 6 devices will show case the performance of OFDMA and DL MU-MIMO operation in the 5GHz band, and lead to a general appreciation of the benefits of using the newer standard. Initial releases of APs supporting the IEEE 802.11ax standard resulted in a mix of Wi-Fi 6 capable and not capable devices in the field. Some expected features like OFDMA were either not present or performed badly. Once a general adoption of Wi-Fi CERTIFIED 6 (per Wi-Fi Alliance standard) APs occur, these early teething issues will disappear. Additional advances in the mixing of Wi-Fi 5 and Wi-Fi 6 devices in 5GHz will occur, as well as more advanced and optimized OFDMA schedulers for Uplink and Downlink handling. It is important for silicon providers to allow scheduler interfaces to be exposed

beyond today's interfaces as there will be requirements and opportunities to customize scheduling of packets and device transmissions based on lots of new innovative ideas, services and solutions.



**Figure 41 - Does 2.4GHz separate out from 5/6GHz in the home over time ?**

There are many options then for the future home Wi-Fi home that could emerge over the next 5+ years

## 8. Challenges and opportunities for Wi-Fi ecosystem

Wi-Fi ecosystem faces significant opportunities with new technologies as was described in previous section. This section describes some of the ecosystem challenges and potential solutions for Wi-Fi.

### 8.1. Enhanced quality of service for applications like VOIP, voice and gaming

Currently QoS implementation, enforcement and management for Wi-Fi is done an individual system level – at client level, router level or at network core level. Client devices like phone and laptops can prioritize traffic in devices using WMM, and DSCP. The implementation and the use of these tags is not uniform end to end. In addition, much of the traffic from devices does not have any priority assigned. Wi-Fi 6 has made significant improvements in prioritizing latency sensitive traffic with support from OFDMA scheduling. However, without accurate tagging from client devices, traffic from applications like VOIP, voice and gaming cannot be prioritized.



A few of the techniques that are proposed to improve QoS are given below.

- Future Wi-Fi standards that can address QoS end to end management

- Automatic classification of QoS for applications using known parameters including UDP/TCP size, frequency, IP address etc.

## **8.2. Onboarding of IoT devices in Wi-Fi networks**

Onboarding of IoT devices into Wi-Fi networks can create potential issues for customers. Many of the Wi-Fi IoT devices including sensors, cameras etc. support only 2.4 GHz spectrum and this creates onboarding issues. IoT devices are onboarded using many techniques including key presses, connecting to Wi-Fi hotspot in mobile device and using other side channel mechanisms. Some mechanisms for device onboarding require both the IoT device and mobile device to be on the same band. Since mobile device may prefer 5 GHz when connecting to router, there are issues onboarding devices that support only 2.4 GHz.

An industry wide solution that encompasses the solutions from WFA including Easy Connect/Device provisioning protocol is needed to simplify customer experience.

## **8.3. Virtualization in Wi-Fi networks**

Virtualization has been used in 3GPP networks starting with 5G networks. Virtualization brings many benefits including simplifying network costs, management of devices and improved network utilization and performance. With the advent of Wi-Fi extenders and mesh networks, virtualization of Wi-Fi networks can yield potential operational improvements for network providers. In multi-dwelling unit deployments, virtualization can simplify operational challenges and reduce the overall network deployment costs by enabling multiple units to share a common infrastructure. In addition, virtualization can be used to provide different quality of services to different classes of devices.


## **8.4. Operational improvements and network monitoring**

Wi-Fi networks tend to have wide variability in quality of network connectivity, primarily due to deployment considerations. Wi-Fi channels may be overloaded due to interference from adjacent networks, use of other unlicensed technologies and due to router placement.

Having robust ability to understand network quality is important for an operator to maintain customer satisfaction with Wi-Fi networks. Network operators have a variety of tools at their disposal including ability to monitor device connection quality, interference in channel, signal strength of device transmissions etc. Some of the newer techniques can also be used that proactively flag network quality issues using anomaly detection techniques.

## **8.5. Devices supporting multiple networks**

Many of the mobile devices support multiple networks including Wi-Fi, LTE/5G, CBRS and public/private versions of Wi-Fi. Traditional mechanisms that primarily default to Wi-Fi networks may not be ideal tradeoff of cost and quality of service. In some cases, operators have multiple networks that can be used by the mobile devices including simultaneous availability of Wi-Fi and LTE/5G/CBRS. Devices may get better network connectivity using LTE/CBRS when outside homes compared to Wi-Fi. A robust mechanism that enables devices to make the right selection is required. This requires collaboration between operators, operating system providers and device manufacturers. This also requires operators to collaborate to create robust mechanisms to share common network infrastructure.



**Figure 42 - Potential synergy between different Operator Wireless networks**

## 9. Conclusion

Since the invention of Wi-Fi and the first packets sent wireless from an AP to a client there has never been a more exciting time for Wi-Fi than the current era of Wi-Fi 6, Wi-Fi 6E and Wi-Fi 7. It is unique in its Wi-Fi cadence and something that needs to be well understood by Cable Operators to plan a 5 year roadmap for their adoption of these technologies and fitting them into their overall customer experience plans. It is clear there will be overlapping SKU's of GW, AP and Extenders for Cable Operators to deploy for different levels of Wi-Fi performance. As Access speeds move to Gbps and beyond then 6GHz becomes the key spectrum to migrate to and to match the Access speeds with LAN performance. However, client upgrades to support 6E and 7 and customer satisfaction of increasing promises of performance will also have to be closely managed. It's clear that Cable Operators can exploit 6GHz immediately by leveraging it for high speed backhaul, cross home wireless backbones and most importantly high performance STB and SMD's as the quality of video increases to 4K and 8K (more driven now by better compression of AV2 and VVC). History has proven that applications will absorb the bandwidth available to leverage and the breach of the 5Gbps, 10Gbps and even 20Gbps over Wi-Fi with scheduled low latency will be quickly filled in by new immersive video and high capacity low latency applications. We typically don't see Cable Operators driving 'Build it and they will come' architectures but there will be fierce battles for consumers connectivity business over the next 5 years with 5G services also trying to leverage new spectrum and drive even Fixed Wireless Connections to overlay Cable Homes. So, the investment in getting 6GHz based services and capabilities into the consumer home to drive new video, immersive services may be the new direction for the Cable Operator and set them up for customer retention for the next 10 years of any new service that emerges. Did I hear someone say Wi-Fi 8?

## Abbreviations

|         |                                               |
|---------|-----------------------------------------------|
| AP      | access point                                  |
| bps     | bits per second                               |
| FEC     | forward error correction                      |
| HFC     | hybrid fiber-coax                             |
| HD      | high definition                               |
| Hz      | hertz                                         |
| ISBE    | International Society of Broadband Experts    |
| SCTE    | Society of Cable Telecommunications Engineers |
| HARQ    | Hybrid Automatic Repeat Request (ARQ)         |
| OFDMA   | Orthogonal Frequency Division Multiple Access |
| SKU     | Stock Keeping Unit                            |
| PON     | Passive Optical Network                       |
| GW      | Broadband Gateway device typically with Wi-Fi |
| STB     | Set Top Box                                   |
| MU-MIMO | Multi User Multiple In Multiple Out           |
| LPI     | Low Power Indoor mode                         |
| VLPI    | Very Low Power Indoor mode                    |
| TWT     | Target Wake Time                              |
| EIRP    | Effective Isotropically Radiated Power        |
| PSD     | Power Spectral Density                        |
| dB      | Decibel                                       |
| dBm     | Decibel-milliwatts                            |
| AV1     | AOMedia Video 1 coding format                 |
| VVC     | Versatile Video Codec                         |
| SMD     | Smart Media Device                            |
| DASH    | Dynamic Adaptive Streaming over HTTP          |

## Bibliography & References

IEEE802.11: *802.11n*, *802.11ac*, *802.11ax* , *802.11be*

# **Augmented Reality Can Improve Wi-Fi Installation In Homes**

A Technical Paper prepared for SCTE•ISBE by

**Ethan Wright**  
Senior Engineer  
Charter Communications  
6360 Fiddlers Green Circle, Denver, CO 80111  
720-378-1502  
Ethan.Wright@charter.com

# Table of Contents

| <b>Title</b>                                                 | <b>Page Number</b> |
|--------------------------------------------------------------|--------------------|
| 1. Introduction .....                                        | 3                  |
| 2. Overview of the Wi-Fi Augmented Reality Application ..... | 4                  |
| 3. Setting up effective Wi-Fi .....                          | 7                  |
| 4. Impact of Augmented Reality .....                         | 9                  |
| 5. Extra Use Cases .....                                     | 10                 |
| 6. Conclusion .....                                          | 12                 |
| Abbreviations .....                                          | 12                 |

## List of Figures

| <b>Title</b>                                         | <b>Page Number</b> |
|------------------------------------------------------|--------------------|
| Figure 1 - APs in different shaped environments..... | 4                  |
| Figure 2 - Edge Detection .....                      | 6                  |
| Figure 3 - Viewfinder and device selection .....     | 7                  |
| Figure 4 - Presence Detection in home .....          | 11                 |

# 1. Introduction

Wi-Fi is a physical medium that is impeded by distance and materials that cause interference. In order to ideally place access points (APs) the physical space needs to be understood. Relevant variables such as spacing, shape of the space, and interfering objects need to be consistently measured. Measuring and evaluating all of these variables creates challenges for technicians to properly optimize networks. APs may be placed optimizing for either cost to the consumer and service provider or coverage. We have built an application that will help a technician or customer design an ideal Wi-Fi space. This paper will outline design issues in Wi-Fi and how an augmented reality application can solve them.


Traditional Wi-Fi design utilizes a meter to collect data about the behavior of the electromagnetic radio waves within a physical space. This provides the technician insight into variables that impact radio waves that are hard to visualize. Interference from other radios and the differing amount of signal degradation caused by materials the radio waves pass through can be interpreted using a Wi-Fi meter. Our AR application combines the traditional data that would be collected using a Wi-Fi meter with additional information about the space such as the area and shape that require coverage.

Additional AP hardware is expensive for service providers and too much coverage overlap can cause negative network performance for clients. If the APs aren't placed close enough together then clients will experience low signal or loss of signal. It is ineffective to train technicians and customers to measure and consider so many variables in order to decide where to place their APs. Instead, service providers may utilize augmented reality to gather data about the physical world and feed that data to systems that may make informed and consistent decisions about how the APs should be placed.

A customer's living space contains three key components of relevant data when deciding where to place APs. The most important factor for effective mesh Wi-Fi design is distance between APs. If the distance is too great dead zones are created and mesh steering becomes ineffective. Conversely, when APs are placed too close to each other interference may be caused on overlapping channels. The service provider also incurs expense due to the additional hardware they are utilizing to cover the space.

The shape of the living space is also an impactful factor when considering AP placement. Wi-Fi radios can cover different shapes based on their design, and location of APs is important to make sure the coverage is complete. For example, if a space has a room adjacent to the square common space, that room may be outside of the coverage for the rest of the living space. As shown in figure 1, non-uniform environment layout can lead to difficult design choices when placing APs. Radios often cover a omnidirectional space uniformly meaning that if a space is not uniform in its shape there will be regions that require a different coverage solution.





**Figure 1 - APs in different shaped environments**

The last relevant variable to be considered is noise and interference caused by foreign materials and appliances. There are some common materials such as mirrors and concrete that can cause a high degree of Wi-Fi attenuation and reflections. Detecting and designing a Wi-Fi system around them is important. The most common device that can cause interference is the microwave, although devices like cordless phones can also cause problems. Locating potential trouble devices and accounting for them can help alleviate problems. This type of interference is hard to visually detect and is often measured through a Wi-Fi meter.

Designing Wi-Fi coverage is not as simple as adding more antennae. It is a problem of design, consideration, and optimization. To design a Wi-Fi system the variables must be known, measured, and weighted. This is a process that is very difficult and not practical for consumers. Technicians are able to be trained as to what variables are important, general guidelines for what thresholds indicate for each variable, and the importance of each. This results in installations that are more efficient than if the customer were to perform the installation, but leads to inconsistencies as much is left to the discretion of the technician. Technicians must spend time measuring and designing the Wi-Fi, which increases the amount of time spent in the home and reduces the number of home installations per day.


## 2. Overview of the Wi-Fi Augmented Reality Application

The augmented reality application is designed to assist with Wi-Fi AP setup by guiding the user to collect all required data. The user is guided through the room and instructed to tap on the borders of the room as well as locating relevant connected devices and APs. Upon opening the application, the user is asked to move through each room, position the camera to view the corners, and then tap on the corner. The user is also prompted to tap on any connected network devices or APs. As the user tags each border or device the augmented reality application will place a virtual tag over the feature, showing the user where the application understands the feature to be. While the user navigates through space the device location and its current RSSI are recorded. After the user has successfully tagged all of the corners and devices they may move their device around the room to see the virtual overlay of how the application understands the environment. This displays to the user the dimensions of their environment and where the variations in signal strength are.

The principles described in this paper have been implemented as a proof of concept application by Charter's Emerging Technology group. The application is an iOS app utilizing ARKit APIs in order to gather relevant data. The application measures the path that the user walks through the home, timestamps


associated with points of the walk, and the relative location of all data points. Once the features have been captured all of the data is sent to a cloud for processing and use by other applications. The APs are also configured to collect high definition RSSI measurements for the AR device during the walk, allowing the device to act as a Wi-Fi meter.

Currently many technologies exist that could make the design of the user interface more accessible for users. The user interface accepts touch input to identify corners defining a room as well as allowing the user to tap on any device in their home that is connected to their network. These events can be seen from an interface perspective in figure 2. This however could be enhanced and even automated using plane detection to automatically locate the intersection between three planes. In an implementation following that pattern, the user would simply move through the space and the layout would be generated as they pass the camera over features. The user could after the fact, or while moving through the space, make any adjustments to correct where the plane detection is erroneous.



**Figure 2 - Edge Detection**

Machine learning can also help automate the placement of network devices. The augmented reality application samples the environment and identifies features including devices. The current implementation allows the user to tap on any device that is currently connected to their network and a list is then provided of currently connected devices for them to select from. This interface is seen in figure 3. Instead of requiring user interaction to select which device is present, image recognition could be employed to match images of known network devices against devices seen in the camera of the augmented reality device.



**Figure 3 - Viewfinder and device selection**

A great benefit of collecting information about the layout of a space and devices within is to enable cloud applications with extensive features. Currently, an application has been developed to consume the input data and render a floor plan for the user and to show their currently connected devices. Many next steps involve the development of new applications using the data, such as an application that would make recommendations for AP placement given the spatial and RSSI variables. Future iterations would also look to expose rules to automation systems like If This Then That or SmartThings in order to provide the ability to interact with their devices based on room or location.

The current working implementation has a lot to expand upon, but even the initial state helps to bring great benefit to the technicians and users. With real data available about the space they are installing APs in, a technician will be able to make a more informed decision about how to design a Wi-Fi system. As the current system doesn't include an engine in the cloud to recommend the location of APs, the technician will follow a process similar to the current process for assessing placement. With the addition of this application technicians will have access to more granular data that maps directly to the criteria they are asked to assess.

### 3. Setting up effective Wi-Fi

Wi-Fi installation is approached differently based on the environment and requirements. Two approaches of relevancy are enterprise installations and retail installations. Enterprise installations are buildings such as restaurants, office spaces, and hotels. Retail installations are environments such as homes and

apartment buildings. The enterprise approach uses an array of tools to gather data about the premises, considers the use cases and number of clients, and utilizes floor plans to add context to the data. In retail installations a technician uses a Wi-Fi meter to map out RSSI in the space and then place an AP if the RSSI falls into the ‘weak’ category. The first approach considers the shape and size of a space and then allows the engineer to design a system based on the data they have gathered. The second approach uses a very simple system to address poor signal allowing for an easily replicable process.

Enterprise grade Wi-Fi installations start with data collection, often move off site for design and architecture, and then conclude with installation and validation. The engineer will either be provided floorplans, or they will create their own by tracing the building. The engineer moves through the building measuring the radio signals and recording them in relation to their location on the floor plan. This is a slow and manual process as the engineer populates the floor plan with all of the measurements. The engineer will also take notes about potential interference locations such as thick walls, concrete structures, etc.

Once all of the data is compiled it will be correlated with information about the use cases the business will experience. This involves understanding the number of clients that will be interacting with each AP and the manner of those interactions. Radios can support different numbers of clients ‘transmitting’ and ‘receiving’ simultaneously. The MIMO configuration allows for fewer or more concurrent clients, and the common radio configurations range from (2x2) MIMO to (8x8) MIMO. In addition to the number of transmitters and receivers, enterprise grade installations need to consider the directionality of the radio. Radios may come as either omni-directional, meaning they cover a 360 degree radius, or directional meaning they cover a longer range in a specific direction. Omni-directional radios are best placed on interior walls and spaced to allow for clients to cleanly be transitioned between APs. Once the engineer has considered all of the measured variables they place their hardware configuration choices onto the floor plan. They return to the site and install the APs based on their architecture documentation. Generally a final sweep is done to verify the design choices and measure the new radio signals.

Retail grade Wi-Fi installations have some shared principles with enterprise, but the process is more rigid and leaves fewer design choices to the technician. In a retail installation a technician is sent to the customer’s home where the Wi-Fi connectivity troubles are being experienced. They will then use a meter to conduct Wi-Fi analysis of the premises and then use a guide to determine whether an AP is necessary. Technicians are asked to rank RSSI in four categories ranging from excellent when greater than -50dBm to weak when less than -70dBm. The technician is asked to determine whether network extenders would remedy a poor network, and in the case that they would they are asked to determine the number of extenders and the location. General rules are provided to the technician in the form of a short video during their training. Technicians are asked to consider the square feet of the building and the shape of the building. No firm guidelines are given, however the technician is informed to distance APs twenty to thirty feet apart when through walls, and thirty to fifty feet apart when in an open space. In their training they are told to consider client usage, density of walls, and to be wary of common objects such as mirrors and household appliances.

Common themes are apparent in both enterprise and consumer grade AP installations. In both situations, the data that is gathered is generally not available after the installation has been completed. This means that there is no easy way of knowing the distance between APs and their efficacy after the fact in both enterprise and retail. The real differences in process are the data that is available to the installer and their level of expertise. Creating a uniform and optimized on boarding experience for enterprise and retail involves making sure the same data points are available to the system that is deciding how to design a Wi-Fi space, and creating a standard method for choosing AP locations.

## 4. Impact of Augmented Reality

In order to optimally place Wi-Fi APs a myriad of data needs to be gathered; The shape and size of the space, RSSI levels at locations within the space, and client use cases are all necessary for effective Wi-Fi design. The augmented reality application asks the user to move through the customer home and move the view over each corner that exists. A user will get within a few feet of the corner and then tap on the corner to record the its location in space. A very similar path needs to be walked to generate a floor plan as is required to sample RSSI. These tools can be combined and provided to the technicians, so they have similar efficacy as enterprise installations. Giving the same tools to technicians that exist in enterprise enriches the data they have access to. Ensuring that the data is stored in a standard way allows the process to get smarter and improved based on past installations.

Augmented reality allows a user to view their environment through a device that overlays relevant metadata into the space. For example, the user is able to point their device at an object, tap on it, and then a virtual tag is added to the object and its location is recorded. The user can visualize all of the recorded tags as virtual stickers they can see when looking through their device at the location.

This interface directions the user to navigate through the space and locate the corners of each room. As corners are identified, a virtual marker is placed over the identified corner and lines are created on the edges. This process gradually builds up a virtual representation of the space the user is navigating through, recording all of the nodes within a relative Cartesian space. These recordings are precise enough to allow for distance and area calculations, meaning that upon completion a record of square footage and shape of the location is fully available as data that can be fed into other systems.

Understanding the physical layout of an environment is critical to effective Wi-Fi design. Measuring the current RSSI helps account for an array of variables that confound proper design, such as interference from appliances and structures. As the user walks through the space to record the relevant corners, they walk a relatively complete path through the environment. As the device moves through the building it's RSSI is recorded by the AP it is connected to. The location measured by the device is correlated to the RSSI recorded by the AP using a synchronized clock allowing the floor plan to be enriched with the data about the radio signals in the space.


In an enterprise system, the architect manually combines the location they take the measurement and RSSI data. They then compare the radio topology with the client use cases, such as number of expected connections. Most modern network installations gather data about what devices are connecting to the network and measure the amount of traffic transmitted and received. These same data points are used by an engineer to decide the proper MIMO configuration for a given space. As the augmented reality application is used to map the environment, network devices may also be identified through image recognition or manual user input. Identified devices may be cross correlated with the devices known to be on the network and then placed into the space. This allows for understanding potential sources of interference, assistance for troubleshooting performance, as well as allowing other applications to integrate with the data. When this data is combined with the data provided by the augmented reality walk, a retail Wi-Fi design may now be as efficient, effective, and fine-tuned as an enterprise installation.

The last variable that differentiates a consumer grade install and an enterprise installation is the knowledge of the installer. The data recorded by the AR application is encapsulated in a standard format and then stored in the cloud. This allows for developers to consume the data after the fact to create models to better advise AP topologies in a space. This removes the potential for human error from the design and allows providers to algorithmically select the variables they would like to optimize for, be it the cost of the hardware or the performance of the Wi-Fi system.

## 5. Extra Use Cases

The data that can be gathered by an augmented reality device is quite useful for Wi-Fi design and optimization of hardware, but its use cases extend further. The application provides relative spatial information for the devices on a network, borders for each room, and distances between all devices and APs. This data can be sent to the cloud and added to the other associated information known about a user's home such as the network devices and known users. Applications can be built to utilize this extra set of metadata to either enrich existing features or curate a new feature set.

One such application is Wi-Fi based presence detection. This is a technology that exploits multipath reflections and other data from the Wi-Fi radio to detect movement between two APs. This allows users “peace of mind” security to be notified when there is activity within their home. One element that is missing from this solution is the difficulty in communicating to the user the location of detected motion. An application such as this could hook into the data provided by the augmented reality application to show the user the room or predicted area where movement was detected between two APs. Currently the platform just simply alerts the user of activity on an unidentified link between two APs. Visible in figure 4 is a home layout with configured APs. This layout shows that a movement between two nodes in this design would indicate movement in a specific room. Providing the user with a graphical representation of their space helps to empower the user and give credibility to applications built on top of the technology.



**Figure 4 - Presence Detection in home**

The Internet of Things (IoT) space can also benefit from the addition of location based metadata to devices and the rooms they exist within. Users are often guided through a cumbersome process of selecting their devices, creating groupings for them, and then naming the groups to indicate where the devices are connected. This is a manual process and is often required to be repeated for every application that contains connected devices. With a service provider able to expose information about what rooms exist in the home and the locations of the devices, IoT applications would be able to automate the forming of groups and creation of rules. IoT applications could now expose rules such as “when a user enters the Living Room, turn on all lights within” and “when movie night is run, dim all lights within ten feet of the television.”

Building metadata about rooms and device location into the dataset for a user is useful for designing a powerful network, yet also provides foundational building blocks for new applications. When enterprise Wi-Fi data is sampled and utilized it’s generally discarded after the topology is designed. Creating a standard system that works in enterprise as well as retail and exposing that data to other applications allows for intelligence to be brought into many feature sets.



## 6. Conclusion

Wi-Fi design is critical to get right when providing service to any customer base. There is a gap between the highest quality Wi-Fi design in enterprise and the methods used to design systems in a consumer's home. This gap may be overcome when augmented reality is used effectively benefiting both enterprise and retail. Proper use of this data enables cloud systems to learn from previous recommendations and grow in capacity to make placement decisions. New systems can be built to help show the customer their environment, communicate where their devices are, and understand their network on a new level. Instead of a user understanding their network as a list of MAC addresses with current DHCP assignments, the user can gain real world context to what is active and where it exists in their home.

Enriching the data that a service provider associates with devices is key to building smarter networks and a broader feature set. In an effort to achieve this, it's important to standardize the approaches taken to design systems and remove as many variables as possible from the system. Instead of giving a technician a set of rules to attempt to apply to a new location, we can gather the data and allow an engine to build the Wi-Fi system. This approach allows the user to better understand their network, empowers the service provider, and gives the technician access to the same tools used by enterprise Wi-Fi architects.

The tools accessible within a mobile device or tablet have evolved to a high degree, it's important to audit new tools as they are available and assess how they can complement service offerings. The processes described in this paper are immediately and directly expected to enhance the capabilities of technicians. A sufficiently intelligent interface should give an end user similar capabilities as a technician in terms of troubleshooting and designing a Wi-Fi network. An external application may make recommendations for where the APs should be placed. The AR application collects and exports all relevant data for designing a network topology. The requirements for running this application are enough compute power to handle augmented reality and a sophisticated enough camera array. This technology has become standard among consumers in the form of their mobile device or tablet. This means that instead of requiring special technology used by technicians, anyone with a modern device can exercise this approach.

## Abbreviations

|      |                                     |
|------|-------------------------------------|
| AP   | access point                        |
| IoT  | Internet of Things                  |
| DHCP | Dynamic Host Configuration Protocol |
| AR   | augmented reality                   |

# **Enabling Automatic Gunshot Detection and First Responders Dispatch for Safer Communities**

A Technical Paper prepared for SCTE•ISBE by

**Wael Guibene**

Director – Wireless R&D  
Charter Communications  
6360 S Fiddlers Green Cir, Englewood, CO 80111  
Wael.Guibene@charter.com

**Hossam Hmimy**

Sr. Director – Wireless R&D  
Charter Communications  
6360 S Fiddlers Green Cir, Englewood, CO 80111  
Hossam.Hmimy@charter.com

# Table of Contents

| Title                                                       | Page Number |
|-------------------------------------------------------------|-------------|
| 1. Introduction.....                                        | 3           |
| 2. AI-based Gunshot Detection.....                          | 3           |
| 2.1. Hardware Platform .....                                | 3           |
| 2.2. Software and AI Platform .....                         | 4           |
| 3. Example Deployment of the Gunshot Detection System ..... | 8           |
| 4. Conclusion.....                                          | 10          |
| 5. Bibliography & References.....                           | 10          |

## List of Figures

| Title                                                                                | Page Number |
|--------------------------------------------------------------------------------------|-------------|
| Figure 1 - HW platform for automatic gunshot detection .....                         | 4           |
| Figure 2 - SW and AI stack for gunshot detection .....                               | 5           |
| Figure 3 - ROC curves comparison SoTA vs. AD.....                                    | 6           |
| Figure 4 - Time domain gunshot detection and audio denoising .....                   | 7           |
| Figure 5 - Example meta-data of gunshot detection and first responder dispatch ..... | 8           |
| Figure 6 - SBC HW platform .....                                                     | 9           |
| Figure 7 - FFMA running AD and AoA/DoA.....                                          | 9           |
| Figure 8 - PTZ camera example .....                                                  | 10          |

# 1. Introduction

Gunshot detection has become a major request from many cities deploying smart city solutions. Current state of the art solutions requires a human interaction mechanism to detect the gunshot incident and number of shots through an operations center (OC). Based on this requirement, current solutions may introduce latencies ranging from 5-10 minutes from the moment a gunshot is detected to alerting of first responders.

According to the FBI, about 70% of active shooter situations are over in under five minutes [1], and the website of National Sheriffs' Association states that "shaving even seconds off the notification and response times can result in vastly different outcomes in these situations." [2]

The response time and latency introduced by existing solutions cannot be reduced as the OC in the loop introduces a human factor that needs to listen to the scene before calling and manually dispatching first responders.

In this paper we introduce a novel gunshot detection mechanism that is fully automated. Our paper describes the software platform along with algorithms that enable:

- Detecting, classifying and localizing gunshots (Audio)
- Recognizing the shooter(s) in the scene (Video)
- Recording the scene from different angles (Video)
- Sending messages in real time to first responders dispatch center as well as authorized personnel with the accurate location and short video recording of the scene.

Our approach applies machine learning (ML) algorithms at the edge to detect and discriminate gunshots from indoor/outdoor ambient sound. When gunshots are detected, the sound direction is also derived, and security cameras follow the direction of the sound to capture the scene, identify the shooter(s) and automatically inform responders about the incident, thus providing significantly improved situational awareness.

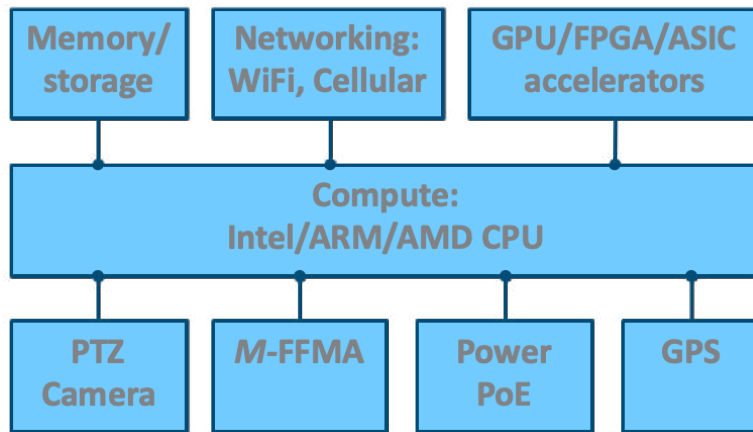
## 2. AI-based Gunshot Detection

### 2.1. Hardware Platform

In our approach, we rely on machine learning and artificial intelligence (AI) algorithms to isolate and identify the gunshots from other sounds.

Our approach consists of hardware (HW) and software (SW) platforms enabling automatic gunshot detection, localization of the active shooting incident and real-time communication to police and first responders.

Figure 1 below depicts the HW platform controlling the cameras.



**Figure 1 - HW platform for automatic gunshot detection**

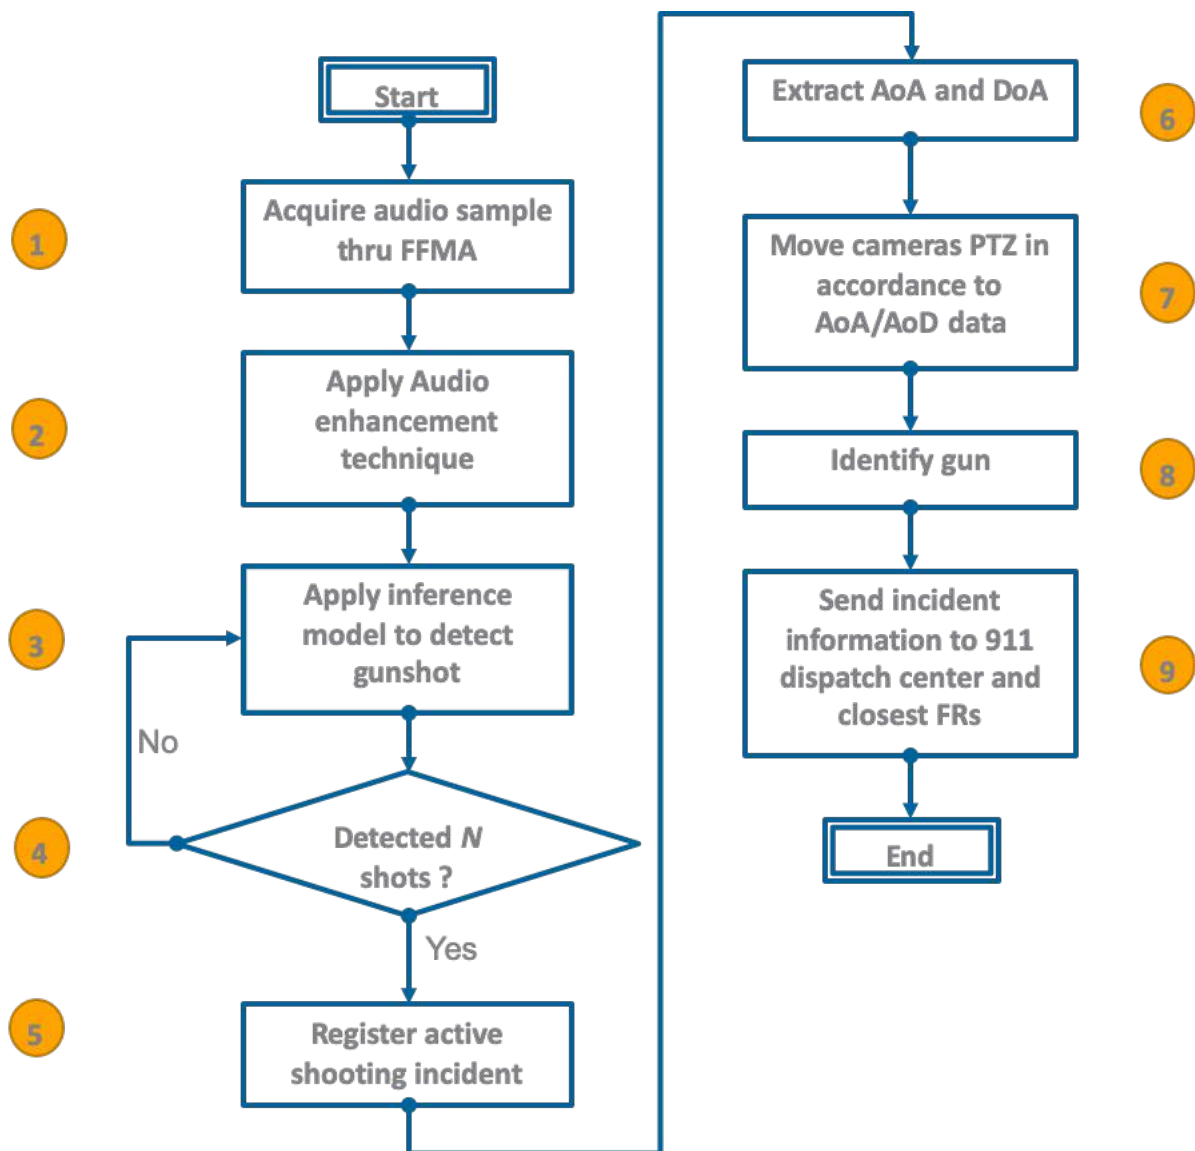
The cameras and their controllers are mounted on street fixtures, e.g. lighting poles or intersections poles.

The HW platform consist of the following components:

- Compute element (central processing unit (CPU) / microprocessor unit (MPU)
- Memory for internal storage
- Networking: either Ethernet, Wi-Fi, cable or cellular for backhauling
- Accelerators enabling different algorithms: graphic processor unit (GPU) for image processing, FPGA for audio and mic-array management, ASIC for image processing acceleration.
- *M*-element far-field microphone array (M-FFMA)
- Pan-Tilt-Zoom (PTZ) Camera: to enable dynamic scene acquisition
- PoE: to enable simultaneous communication and power delivery
- Global Positioning System (GPS): to localize the HW platform and fine timestamp events of active shooting

## 2.2. Software and AI Platform


Figure 2 below describes the end-to-end flow of the automated gunshot detection and reporting framework.



**Figure 2 - SW and AI stack for gunshot detection**

1. Audio samples are acquired through the M-element far field microphone array (FFMA). The FFMA allows active noise cancellation (ANC) and direction of arrival (DoA)/angle of arrival (AoA) algorithms to run simultaneously with high accuracy.
2. Given that the gunshots have spike-like behavior in the audio temporal domain, we introduce the algebraic detector (AD) to act as both a detector and an ANC for the audio captured by the microphone array. This will increase the accuracy of the ML inference model that detects the shots. The AD is a pre-processing step that locates the shot in time domain prior to applying our ML model. This will reduce the  $P_{\text{False Alarm}}$  for the end-to-end (E2E) application.

One of the metrics that confirms the accuracy of the AD as a joint ANC and denoising technique is highlighted in the receiver operating characteristic (ROC) curve in figure 4 below and the example of gunshot detections in time domain. From the ROC curves, up to 70% accuracy at 1% false alarm rates are achieved using just 30 samples. This will scale to over 95% accuracy with a constant probability of false alarms at less than 1% at full MPEG Audio Layer-3 (MP3) audio quality (44100 samples).




**Figure 3 - ROC curves comparison SoTA vs. AD**

We can see from the Figure above that the ROC curve associated with AD outperforms both NEO and Wavelet detection algorithms.

Another metric that helps assess the AD performance is an example of gunshot detection in time domain and the capability of the AD to denoise and separate the spike-like behavior of the gunshots from other noise at a location.

The Figure below shows an example of ground truth and isolated gunshots from city noise. The AD can run either offline or in real time to simultaneously reduce noise and locate spike-like shots in time domain.



**Figure 4 - Time domain gunshot detection and audio denoising**

3. After detecting the likelihood timing of the shot, we apply a pre-trained model on gunshots over the cleaned audio signal. The pre-processing ANC through AD step allows for cleaner, better audio signal to feed into the ML classifier and allows better results with fewer false alarms for the E2E application.
4. If a gunshot is detected, a counter and a timer turn on. The counter is set to a value of N shots and the timer is set to P seconds. We set these conditions to further ensure robustness of the solution. As we do not trigger an alert at each shot, we need to aggregate N shots within P seconds to classify the event as an active shooting incident and trigger the rest of the algorithm. Shots fired from different people are processed concurrently.
5. If the algorithms register N shots within P seconds, it triggers an active shooting event. The algorithm will keep a counter increment of each shot it is detecting in the background.
6. Using the FFMA capabilities to run AoA and DoA, the system infers the angle and direction from which the shots are being fired.
7. The system issues pan-tilt-zoom (PTZ) “OnVif” commands to the camera to pan and tilt to the direction of the active shooting. The zoom is calibrated thru the Received Signal Strength Indicator (RSSI) of the audio track to estimate how far the shooting is from the camera and the microphone array. Depending on the deployment settings, the gunshot is most probably detected by more than one system. Each system is commissioned with its GPS coordinates and the shot is either identified by the closest detection system (if there is only one system deployed) or by triangulation of the different systems. The system starts recording and storing the scene in the camera’s internal storage as well as exposes the live stream thru real time streaming protocol (RTSP). The system also runs on-camera facial recognition and captures all faces in scene.
8. Another advantage of our approach is the capability to recognize guns by correlating the noise-free sound provided by the AD and ANC to an on-device database containing the sound signature of the most sold guns.
9. A standardized report is created by the system and sent to first responders. An example of the aggregated data sent by the system is as described by the following JavaScript Object Notation (JSON) meta-data, in Figure 5, with the following fields:
  - **NB\_Shots\_Per\_Platform**: identifies each HW platform that detected the shooting as well as the number of shots detected per platform.
  - **AS\_GPS\_Coordinates**: either the trilateration GPS location of the shooting or the closest HW platform that detected the active shooting.
  - **Timestamp**: timestamp at which the alert was sent.
  - **Video\_Streams**: links to the live-feeds from each HW platform that detected the active shooting for first responder and 911 network operations center (NOC) to evaluate in



- **Faces\_Captured\_At\_Scene:** as the camera is capturing the scene, we run in real time face detection of all people present at the scene and store and send images as part of the JSON in base64 encoding.
- **Gun\_Type:** is the identified gun involved in the shooting alongside its presence likelihood

**Figure 5 - Example meta-data of gunshot detection and first responder dispatch**

In this section, we give concrete examples and results from a proof of concept (PoC) that is running the gunshot detection system indoors in an office setting.

© 2020 SCTE•ISBE and NCTA. All rights reserved.



**Figure 6 - SBC HW platform**



**Figure 7 - FFMA running AD and AoA/DoA**

Figure (8) shows one of the PTZ cameras used that is connected and controlled by the automatic gunshot detection system.



**Figure 8 - PTZ camera example**

## **4. Conclusion**

In this paper, we presented a unique and novel approach to an AI-assisted automatic gunshot detection and reporting system. Our approach enables fully-autonomous isolation of gunshot sounds and recognition of gun types as well.

We have deployed and tested our solution in our lab using recorded gunshots to test the resiliency of the algorithms and test the end-to-end performance and latency of the solution. Our tests show conclusive results on the detection time: near real-time detection with under two seconds to detect an active shooting situation, take photos of the likely shooter and the scene and send them to first responders and police for dispatch.

Future work includes testing the system using real gun shots in the shooting range, as well as collecting more gun sound samples from the shooting range to enhance the model.

## **5. Bibliography & References**

[1] ShotSpotter: Home [<https://www.shotspotter.com/>]

[2] Blair, J. Pete, and Schweit, Katherine W. A Study of Active Shooter Incidents, 2000-2013. Texas State University and the Federal Bureau of Investigation, U.S. Department of Justice, Washington D.C., 2014.

[3] “Embracing Technology to Decrease Law Enforcement Response Time”, Feb. 18, 2016. National Sheriffs’ Association ([www.sheriffs.org](http://www.sheriffs.org)).

# **Repair The Ides Of March:**

## **COVID-19 Induced Adaption of Access Network Strategies**

A Technical Paper prepared for SCTE•ISBE by

**Dr. Robert Howald**

Fellow

Comcast

1800 Arch Street, Phila, PA 19103

215 286 8037

robert\_howald@comcast.com

# Table of Contents

| Title                                                      | Page Number |
|------------------------------------------------------------|-------------|
| 1.0 Introduction .....                                     | 4           |
| 2.0 By the Numbers .....                                   | 4           |
| 2.1 By the Geography .....                                 | 8           |
| 3.0 Online During a Pandemic: What Are We Doing? .....     | 9           |
| 4.0 Post-Pandemic: So...What Might Normal Look Like? ..... | 12          |
| 5.0 Next Generation Nimble .....                           | 13          |
| 5.1 The Near Future Has Arrived Ahead of Schedule .....    | 14          |
| 5.1.1 It's the Bandwidth, Stupid .....                     | 15          |
| 5.1.1.1 Rich Fiber Diet .....                              | 16          |
| 5.1.1.2 Widen the Lanes .....                              | 17          |
| 5.1.1.3 Every Bit Counts .....                             | 18          |
| 5.1.2 Be There for Me Always .....                         | 21          |
| 5.1.3 Be Snappy .....                                      | 26          |
| 6.0 Conclusion .....                                       | 29          |
| Abbreviations .....                                        | 31          |
| Bibliography & References .....                            | 32          |
| Acknowledgments .....                                      | 33          |

# List of Figures

| Title                                                                                         | Page Number |
|-----------------------------------------------------------------------------------------------|-------------|
| Figure 1 - COVID-19 Cases vs Time in the United States [15] .....                             | 5           |
| Figure 2 - COVID-19 Induced Downstream Traffic Growth [14] .....                              | 6           |
| Figure 3 - COVID-19 Induced Upstream Traffic Growth [14] .....                                | 6           |
| Figure 4 - Regional Traffic Variation are Significant [14] .....                              | 8           |
| Figure 5 - COVID-19 Impact on Video Conference and Chat Applications [9] .....                | 9           |
| Figure 6 - COVID-19 Impact on Weekly Video Conference Traffic [2] .....                       | 9           |
| Figure 7 - COVID-19 Impact on Peak Busy Hour Time Shift of Upstream Traffic [1] .....         | 10          |
| Figure 8 - COVID-19 Impact on Weekly Virtual Private Network Traffic [2] .....                | 10          |
| Figure 9 - COVID-19 Impact on Streaming Video Traffic [2] .....                               | 11          |
| Figure 10 - COVID-19 Impact: Web Access vs Mobile Access [10] .....                           | 11          |
| Figure 11 - COVID-19 Impact on Gaming Traffic [2] .....                                       | 12          |
| Figure 12 - N+0 Architecture Basics: No RF Amplifiers and Expanded DS and US BW .....         | 16          |
| Figure 13 - Fiber Deep Cost Effectiveness Compared to Continued Node Splitting Only [6] ..... | 17          |
| Figure 14 - Upstream Lifespan Expansion Options [1] .....                                     | 18          |
| Figure 15 - Multiple Modulation Profile Potential of D3.1 .....                               | 19          |
| Figure 16 - DOCSIS 3.1 Multiple Modulation Profiles vs Time and Freq [8] .....                | 20          |
| Figure 17 - Programmable Modulation Application Workflow [12] .....                           | 20          |
| Figure 18 - Exploiting More DOCSIS 3.0 Upstream in 5-42 MHz Systems .....                     | 21          |
| Figure 19 - MER Distribution for a Large Sample of CMs on a N+0 System .....                  | 22          |
| Figure 20 - Reduction of Trouble Calls in Fiber Deep Deployments .....                        | 23          |

|                                                                                                                                                             |    |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| Figure 21 - Real-time Dashboard of Network and Component Availability and Uptime.....                                                                       | 24 |
| Figure 22 - The Virtual Cloud Efficiently Shifts Capacity Where it is Needed.....                                                                           | 25 |
| Figure 23 - Sources of Latency in DOCSIS 3.1 Networks [13] .....                                                                                            | 26 |
| Figure 24 - Traditional and LLD Service Flow Architecture [13] .....                                                                                        | 27 |
| Figure 25 - Industry LLD Efforts are Promising for Supporting Real Time Services Such as the Recent<br>Explosion of Video Conferencing Due to WFH [13]..... | 28 |

## List of Tables

| <b>Title</b>                                                                                | <b>Page Number</b> |
|---------------------------------------------------------------------------------------------|--------------------|
| Table 1 - Capacity Management Dashboards Guide Targeted Network Upgrade Investment [1]..... | 16                 |

## 1.0 Introduction

The instant network stress test that resulted from COVID-19 is something the most precognizant of network engineers could not have foreseen. Even if they had, the most convincing of them would have little chance of swaying financial teams to increase investment for a traffic spike to come, and a ramp to all-time high sustained usage plateaus that would come in a matter of days.

Most importantly, the performance of the network during this most demanding and critical period shone a light on the industry in a way no press release, marketing campaign, or new service feature could. Nothing was more essential to our customers than delivering their known services, on steroids, robustly and 24/7.

This paper will review the COVID-19 period in five parts. First, we look at the raw numbers as work-from-home, school-from-home, zoom-from-home, and virtual happy hour became everyday behaviors. We will review the data and applications, downstream and upstream, against network capabilities and explain why things went well and where there were challenges.

We will then look at the response from three perspectives. First, we will assess utilization metrics in the most challenged areas, evaluate vulnerabilities, and describe quick response actions that maintained network health.

Secondly, we discuss acceleration of near term initiatives to further enhance capacity in these areas, but also more broadly to address the national rising traffic tide.

Lastly, we will zoom out and discuss how the dynamics of 2020 effect the strategic network plan.


In the final part of the paper, we will discuss the “new normal” of usage, applications, and behaviors, and describe how network planners should consider them in strategy development. Among the silver linings of the pandemic is confirmation that the industry has been focusing on, and executing on, all of the right things over the years leading up to it, and the network is indeed prepared for scenarios that are, literally, beyond imagination.

## 2.0 By the Numbers.....

The two months after the first reported case of COVID-19 in the United States (January 21, 2020) dispelled the notion that the country might escape catastrophic ravages of a pandemic that began overseas, the dreaded “community spread” of the virus kicked into gear. In a period of time on the order of weeks, a population that had been largely not focused on COVID-19 was experiencing some of the most extreme changes to life-as-we-know-it in a generation.

Just as incredibly, an initially bewildered population adapted to this new normal, following rules and recommendations from government leaders for what would otherwise be considered draconian civil liberties restrictions. Furthermore, the rules and recommendations changed frequently and varied from state to state. Nonetheless, the US population, at least initially, rallied behind “we’re all in this together” to achieve the objective of flattening the curve. Indeed, all “you-know-what” broke loose later in the Spring, but that’s left for another genre of publication to discuss.

Figure 1 captures this COVID-19 case-tracking period, showing the rapid rise that began in the second half of March.



The 7-Day moving average of new cases (current day + 6 preceding days / 7) was calculated to smooth expected variations in daily counts.


**Figure 1 - COVID-19 Cases vs Time in the United States [15]**

As admirable as the American people may be, one of the major contributors to the adaptability and relative smoothness of the transition into this unusual new routine was the capability and performance of the broadband network. It allowed for the virtual connectivity essential for the transition, and enabled high volume use of bandwidth-heavy, video-rich applications in homes to be exercised at unprecedented scale, often to substitute for activities no longer available to individuals under “shelter in place” orders. In addition, while massive furloughs and increased unemployment are tragedies caused by the pandemic, the broadband network provided a means for a large percentage of the population to work from home and do so effectively. It is estimated that about 29% of the US workforce is able to effectively work from home [11]. The network also enabled schools to continue to operate, albeit at a reduced level of intensity and effectiveness.


Figure 2 and Figure 3 show the Downstream (DS) and Upstream (US) traffic growth, respectively, over time, as shelter-in-place struck. Most honest broadband technologists would tell you that if you had predicted on March 1 that their networks would see traffic spikes over a period of two weeks like those shown in Figure 2, they would have expected some crash-and-burns in the field, and accompanying negative media coverage.

Rather than a spike of Internet outages, however, what was observed was that while there was an occasional connectivity “housefire” (which was able to be extinguished relatively quickly), in fact there were simply not huge numbers to address. Generally speaking, the impact was absorbed well due to years of experience of managing network traffic, understanding peak and average utilization thresholds that effect the customer experience, and continuously preparing and optimizing the network for traffic growth [1]. Internet traffic growth has not abated since the launch of cable broadband in the late 1990s; only the pace of growth changes year to year – the so-called Compounded Annual Growth Rate, or “CAGR.” Network engineers responsible for traffic management of their systems speak and think CAGR in their sleep, with an occasional nightmare about 8k streaming holograms.





**Figure 2 - COVID-19 Induced Downstream Traffic Growth [14]**



**Figure 3 - COVID-19 Induced Upstream Traffic Growth [14]**

Let's look at the COVID-19 spike in the context of CAGR. DS and US CAGR vary, of course, and historically the DS has grown 40-50% per year, while the US has been closer to 20-25% per year and has been more dynamic year-to-year. Upstream years of 15% have been recorded as well as very aggressive years, such as when peer-to-peer file sharing took off in the early 2000s. More recently, home security cameras streaming video are becoming noticeable contributors as their popularity increases, and this increase is tied to smart home automation trends in general.

In recent years for the Downstream, many over-the-top (OTT) video service options have fully matured into the marketplace. And yet, people have not grown more than two eyeballs to consume the

voluminous programming. It is a bit early to make the call of a downward trending DS CAGR, but it has settled into the 30% range.

Below are some very useful “napkin math” rules of thumb (pre-COVID speak alert – oh, how great will it be to be able to grab a bar napkin to do math, diagram architectures, and sign contracts!). We can relate CAGR to the so-called Traffic Doubling Period (TDP):

TDP = 2 years when CAGR is ~40%

TDP = 3 years when CAGR is ~ 25%

TDP = 4 years when CAGR is ~ 20%

Now, consider the COVID-19 traffic spikes in Figure 2 and Figure 3. Downstream growth peaks at about 20% and has since receded to 10%. This tapering to 10% coincides with the loosening of restrictions, but also with the onset of spring and summer weather. This year more so than any other people were anxious at the first opportunity to get outdoors. Furthermore, the summer broadband season is typically slower. So, it may be premature to assume that the 10% is a settling point for DS CAGR after the pandemic. If we use the more conservative guideline (from an architect’s perspective) of 20%, at least until we know more about what a post-COVID new normal looks like, then the virus has accelerated Downstream traffic growth by 7-8 months, using 35% as a reference DS CAGR. The Compounded Monthly Growth Rate (CMGR – unfortunately not as easily expressible as “CAGR” in conversation!) for 35% works out to 2.5%. It seems clear that some behavior and societal impacts will remain after the virus is no longer a threat large enough to warrant restrictions.

For the Upstream, an initial spike of 35% has receded to about 25% as of July of 2020. Using an US CAGR of 25% means that traffic has accelerated by just about one year. The 35% spike, however, if sustained, is an acceleration of about 16 months (1.9% per CMGR). Generally, cable operators are driven to action in the network by congested upstream. The primary reason for this is simply that there is much less upstream, home-to-headend spectrum available, and it is more difficult to use. The network was built as a one-way video broadcast network originally, so Downstream spectrum was emphasized. Fortunately, analog video is so sensitive to noise and intermodulation distortion that the network was built with very high fidelity, even after being carried through dozens of amplifiers, in the earliest days of cable, enabling today the opportunity with fiber and DAA to have tremendous digital communications efficiency. So, the limited upstream has been able to be made quite powerful, given the available spectrum, which is typically 5-42 MHz in North America – or about 5% of total available capacity, depending on where you draw the upper spectral boundary for the downstream signal path. Cable operators have dealt with this limitation of available US spectrum since the launch of High Speed Data services (HSD), in the late ‘90s, and of course benefit from the fact that the residential asymmetry ratio of DS to US traffic is about 12:1 today, and has never ventured below about 7:1.

The bottom line is that while the US has been exploited very effectively, it more often than not drives plant upgrade strategies. Several access network initiatives in the industry are about adding US spectrum via 85 MHz Mid-Split, 204 MHz High-Split, and DOCSIS 4.0 Full Duplex and Extended Spectrum. The COVID-19 acceleration of 16 months is especially meaningful in terms of its impact to US ports close to the congestion threshold. A traffic spike such as this would normally occur over many months, and under normal circumstances that yields sufficient time to plan the upgrades necessary to avoid oversubscribing the link.

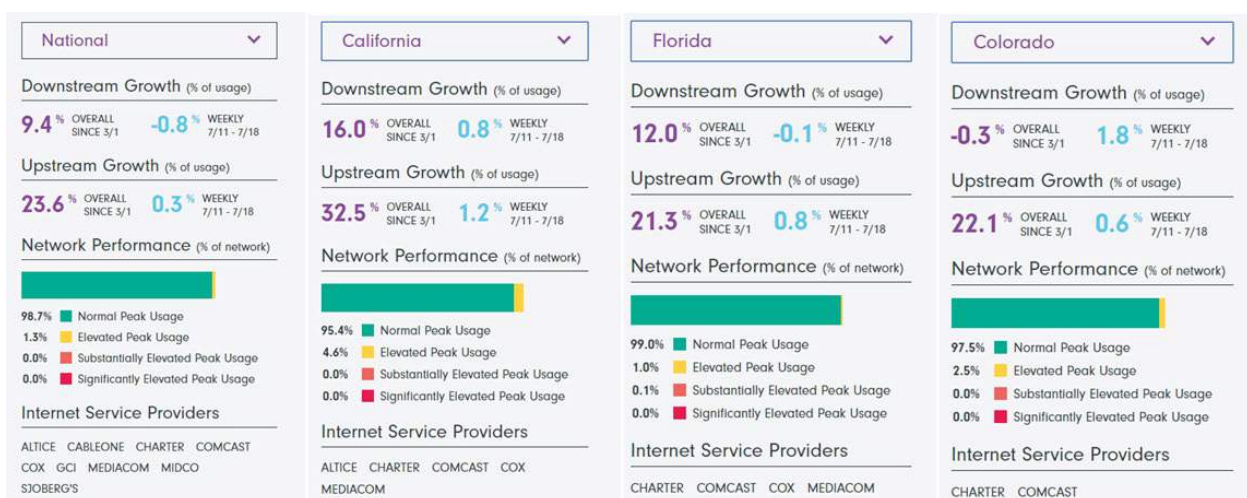
On the flip side, operator awareness of US bottlenecks had already been quite high, so the tools and strategies to deal with it were well known and in place to manage it. The difference with COVID-19 is

that the multi-year plans lost the luxury of much of that prep time and budgets, when it accelerated forward. It is in scenarios where fiber nodes were on or approaching utilization alarms within a year that network congestion was triggered – the aforementioned “house fires” – where work to resolve it needed to begin immediately.

## 2.1 By the Geography....

Averages provide a general understanding of network behaviors and trends. However, they can be dangerous as well, as regional traffic variations and network performance are tied to the issue that matters most – the customer experience. Where technology-centric businesses, professionals, and universities are key tenants of a region’s environment, populace and demographic, average and peak traffic metrics and growth will be correspondingly higher than in locations without these characteristics. So, as valuable as National averages are, operators have to be much more surgical in their knowledge, and in their response. Regional variations are too significant to have a one-size-fits-all network strategy, under normal circumstances, much less considering something as dynamic and unprecedented as a pandemic spike.

Figure 4 is a snapshot the network numbers for several representative states, all of which are served by Comcast as well as other cable operators.




**Figure 4 - Regional Traffic Variation are Significant [14]**

What is clear from Figure 4 is that the National average would be insufficient as a guide for a COVID-19 response plan for California but might hit the mark in Florida. Alternatively, Colorado appears to be an area where prioritizing US should definitely take precedence over the DS. It is often the case that when addressing congestion in one direction, such as upstream, it is cost effective to address the other direction at the same time.

A conclusion that can be drawn, and a reality of the experience in the early weeks of the pandemic, is that reactive operations to put out the “housefires” were not uniformly distributed. High tech corridors, such as exist in California and other parts of the country, accounted for a disproportionate share of rapid response efforts. Pandemic or no pandemic, non-uniform implementation or staggered timing of network strategies is a common component of access evolution.


### 3.0 Online During a Pandemic: What Are We Doing?

It is clear and by now intuitive what applications fueled the surge in traffic, including why the US has seen the larger increase. Figure 5 shows the rapid rise of video conference and video chat applications in response to the shelter-in-place lockdowns. Such a rapid spike over a short period of time is exactly the kind of thing that a network engineer loses sleep over. The increasingly popular business conferencing tool Microsoft Teams is overshadowed by the Zoom application, which exploded as the pandemic hit, as a physically-distanced, family-friendly chat experience. However, in relative growth terms, Teams also exploded exponentially, going from 560 Million minutes on March 12 to 2.7 Billion minutes by March 31, or about 70% growth per week during that period of time [4].




**Figure 5 - COVID-19 Impact on Video Conference and Chat Applications [9]**

Figure 6 shows the traffic on a weekly timeline, highlighting how much can be attributed to work-from-home (WFH) employees, as evidenced by the dip in traffic over the weekend.




**Figure 6 - COVID-19 Impact on Weekly Video Conference Traffic [2]**

To punctuate this conclusion, Figure 7 shows how the upstream has experienced a shift in its Peak Busy Hour (pbh) usage, from about 9 pm local time, to 2 pm local time. The biggest percentage increase over pre-COVID utilization is actually at 1 pm, but because the characteristic of the US is a gradual rise throughout the day, the absolute peak has moved to about 2 pm [1].




**Figure 7 - COVID-19 Impact on Peak Busy Hour Time Shift of Upstream Traffic [1]**

Finally, Figure 8 shows the predictable spike of virtual private network (VPN) traffic at the onset of the lockdown, such as would be expected when a massive shift to WFH has taken place.



**Figure 8 - COVID-19 Impact on Weekly Virtual Private Network Traffic [2]**


Moving onto other expected online implications of shelter-in-place, Figure 9 shows the increase in use of streaming video usage, as alternative entertainment options – movies, concerts, etc. – were eliminated. The taper not long after the upward trend is attributed, again, to a combination of spring weather providing opportunities to relieve cabin-fever, and possibly, “binge watching fatigue,” as reality set in following initial estimates of short term lockdowns. Note that the call-out labeling on this figure appears to be in error – the highlighted dates appear to be weekends.



**Figure 9 - COVID-19 Impact on Streaming Video Traffic [2]**

Not captured in any of the charts but also intuitive is that the applications identified have seen growth in being accessed through web URLs instead of mobile applications. Figure 10 shows how trends for popular applications representing streaming video shifted rapidly away from apps and towards websites. The relevance to this paper is somewhat secondary. The large increase in streaming video traffic, which makes up approximately 70% of all web traffic, is still the main story. However, the secondary effects are twofold:

- 1) Applications accessed on small screen devices will on average translate to lower bit rates of video compared to their desktop or larger screen/tablet counterparts
- 2) Web access as well as the mobile devices themselves are more often accessing content via Wi-Fi instead of mobile networks, driving more traffic to the wireline broadband network.



**Figure 10 - COVID-19 Impact: Web Access vs Mobile Access [10]**

A final interesting traffic dynamic is shown in Figure 11, which captures gaming application growth. Again, with entertainment options limited, and gaming already a popular and growing online activity, regular gamers had less reason to back away from the console, and casual gamers had good reason to increase usage. An interesting observation is that, while Figure 9 has the familiar peak and lull pattern



common to weekly residential traffic behavior, the pattern itself was altered in the gaming space, at least in the early weeks. Weekdays and weekends blended into one continuous gaming binge! It will be interesting to see over the long term if gaming was a temporary diversion or if it also has achieved a new baseline level of significance and permanence due to once-casual gamers becoming more hardcore participants.



**Figure 11 - COVID-19 Impact on Gaming Traffic [2]**

## 4.0 Post-Pandemic: So....What Might Normal Look Like?

As mentioned at the outset of this section, in retrospect we have a largely intuitive outcome from the pandemic restrictions. While it is difficult for anyone to “crystal ball” in such an unprecedented situation, there are some consensus beliefs emerging about the Post-COVID-19 “new normal” with respect to traffic and the network. Permanent changes of the traffic baseline above the pre-pandemic steady state is an expected outcome. Summarizing some of this emerging conventional wisdom:

- It will be truly awesome when we will be in a position to compare projections with the real Post-COVID-19 traffic data!
- *Working from Home (WFH)* – WFH will play a larger role for office workers. The comparable productivity observed, for companies where this was the case, will have companies reconsidering the potential savings associated with the operating expenses of commercial office complexes. People density in an office building may be changed forever, and trends like open floor plans rethought. Video conferencing had been adapted at Comcast by the end of 2019, and in many other companies as well. With an expected permanence to more WFH, applications like Microsoft Teams, WebEx, Zoom, etc., which have been indispensable from the outset of the pandemic, will be more essential. Also, their capabilities are likely to improve with a higher focus on scale, emulating face-to-face features, and blending in attributes that f2f meetings cannot obtain, such as the virtual backgrounds (I’m in Hawaii! I’m on a ski lift!) and associated computer-aided graphics seen in some of these applications already today.
- *Normalization of Online Education* – Universities have been offering online classes for many years, and some universities are entirely online. However, for most full-time universities and

certainly primary and secondary schools, face-to-face instruction is still considered as the most effective means of instruction. Online courses offer much more convenience, if a somewhat less effective alternatives. However, tools focused on a classroom experience will now get significantly more attention in a bid to emulate true classroom environments, while offering capabilities only a digital environment can easily do, such as shared whiteboarding, offline non-disruptive chat, and an audio and video record of the experience.

- *Normalization of Virtual Social Gatherings* – The “virtual happy hour” has become a ritual among my spouse and her friends and with family. Despite no ritual clinking of glasses to “cheers”, shelter-in-place orders caused a much wider interest in applications like Zoom. A positive side effect of requiring such applications for any social contact is that through this new familiarity, families who were physically separated by virtue of living in different parts of the country for years began to use it to virtually visit each other far more often than they otherwise would have. In this way, the pandemic has brought families which are naturally distanced from one another into more frequent contact. It remains to be seen if that continues post-COVID, but it seems likely that the impact can only be a net increase in these applications.
- *Preparation for the Next Apocalypse* – Many businesses reliant on personnel in factories, warehouses, logistics centers, etc. suffered substantially with the absence of these employees, especially in those industries not considered “essential.” Production, logistics, supply chains, employment itself, etc. were disrupted, magnifying economic woes. Companies with bread and butter in goods and services tied to manufacturing, shipping, and logistics are expected to invest in increased digitization and automation of operations, including use of more machine-to-machine communications (M2M) as part of robust contingency planning, to limit the financial damage that future pandemics or unforeseen interruptions of operations can cause.

An interesting pandemic perspective, from a sales point of view is that, as a salesperson, face-to-face meetings, lunches, dinners, etc. are a key element of the business relationship. Sales teams are generally in a similar boat regarding access to face-to-face meetings. While established accounts are generally as vulnerable (or not) as always, given this equal playing field, new business accounts are increasingly difficult to land, as the reliance on relationship building and face-to-face discussions are magnified.

If this is so, the pandemic effect may be detrimental to disruptors, as well as to new business areas in the incubation phases for established OEMs. Beyond the survival mode drag placed on technology development already, this could result in a stall to innovation in the industry as a whole. As technologists and representatives of the industry with a mutual desire to keep our ecosystems healthy, this is something to keep an eye on. There has been an appropriate focus on “keeping the lights on” activities and minimizing risk to keep people at least virtually connected. However, given traffic acceleration impacts of COVID-19, it is even more important that innovation does not suffer a setback as technologies to deal with it have been essential.

## 5.0 Next Generation Nimble

The above data tells a clear story and suggests an expected post-pandemic permanence to some of these trends. That said, the job of any access network team is to determine what revisions may be necessary for the path forward. We will note that the broadband network is but one of many potential traffic bottlenecks and addressing it in isolation will not necessarily eliminate all customer issues. Many of the mature applications identified within, that are shifting with the trends of the pandemic were, like the



network, operating in a “business as usual” (BAU) pace of upgrades of features, capabilities, and scalability. The application servers themselves, and interfaces to and from them, also experienced an unforeseen step function increase in demand.

VPN servers, for example, perhaps accessed routinely by 15% of a company’s remote workforce --with compute and I/O sized accordingly -- suddenly saw perhaps 4-5x the number of remote employees vying for access, slowing everyone down. A video conferencing application, for example, may be fully suitable for cross-country meetings from several remote offices, but ill-suited for 30 individual streams. Lastly, we have seen in the news the challenges the Zoom application has had with security of sessions, a feature important to all, but with very significant implications to business data compared to the virtual family video chat. So, as mentioned, while network engineers revisit their strategies and plan for evolving the network, there are many parallel efforts across the products and applications such as these, and the companies who own or operate them. The end performance for customers involves successful migration and execution on both fronts.

Lastly, we note that the focus will be on the edge and access network, which is generally the cable system bottleneck. But we should note that the in-home broadband network, in particular the multiple and simultaneous user Wi-Fi experience, may also come under duress as a bottleneck. The new generation of Wi-Fi 802.11ax (Wi-Fi 6), and the new 6 GHz band are arriving at just the right time to improve the in-home experience. In-home Wi-Fi mesh extenders have also matured to provide a significantly improved coverage experience. Depending on their current Internet service and gateways, customers may need to upgrade their in-home network.

## **5.1 The Near Future Has Arrived Ahead of Schedule**

We have seen the drivers of new behaviors, and, to the extent that we can use the word “predictable” in a discussion based on an unfinished global pandemic, the network traffic aspects and key applications of pandemic life are, in fact, predictable – except perhaps that gamers evidently never EVER take breaks. Furthermore, we can recognize that the pandemic has not changed very much materially about what the important initiatives for advancing the access network are. All of them are things that we should have been thinking about already.

What has changed, however, is the execution timeline of these strategies, including the relative priorities and the proper sequence of initiatives. And, importantly, it has increased greatly the level of awareness of our customers, and the nation as a whole, to the criticalness of the network and availability. This plays a role in the approach to phasing, execution, and risk assessment, as initiatives are brought to trials and rolled out. In access network engineering, it has normally been the case that the job is well done when major network-impacting technology upgrades take place and the customer does not notice you were there, unless it’s having a new product offering to consider, or because they cannot remember the last time they had to call to talk to a care agent on the phone. As the shelter-in-place orders came out and companies started dispatching their employees to work from home, ALL eyes were on the network. And broadband engineers everywhere should be very proud of what they built and how this situation has played out so far.

Armed with the trend data shown earlier in this paper, the recognition that the trends are likely not temporary, that precise accuracy of end-state traffic behavior will be uncertain for some time – yet the ability to surgically address rapidly changing dynamics has never been more essential – a recalibration of the network evolution makes sense. This should be done with the following considerations in mind:

1. *Enable Agile Coaxial Capacity* – Quickly add capacity, above and beyond “BAU” increases, necessary to recover the pre-COVID bandwidth-versus-time runway, and in anticipation of increasing US CAGR and that the observed trends so far get rooted beyond the first wave of adopters.
2. *Hardening and Resiliency* – Enhance the reliability and availability to enable a seamless WFH and online education experience for all types of remote workers and students. Increase visibility and proactivity to the customer experience. Create a resilient and self-healing network.
3. *Responsivity* – Drive down network-induced latency and jitter to support the increase in real-time applications, such as video conferencing, and to ensure the increased in gaming traffic does not impact the gamer experience.
4. *Enabling of Enterprise Services* – Perhaps this seems out-of-touch with the obvious softness in office-based business services. However, economic losses suffered due to shutdowns of industries dependent on “non-essential” workers are expected to lead a drive to digitization of processes and industry. This includes M2M communications, coupled with data analytics, in order to be less vulnerable in the future.

### 5.1.1 It's the Bandwidth, Stupid

Bandwidth, bandwidth, bandwidth. When is it ever NOT about the bandwidth when it comes to access network evolution? Well, this is NOT the exception.

One of the valuable lessons confirmed by the pandemic is the value of building ahead of the traffic curve. How far ahead is an ongoing debate and a target that shifts around many technical, business, competitive, market, financial dependencies. Generally it is good to have several years of network runway ahead at all times, simply because we KNOW traffic will continue to grow and because it is only every so often that the opportunity to touch the network takes place, and certainly not continuously. It is this built-ahead-of-the-curve discipline that allowed network operators to not just survive the pandemic storm but excel in it. The spike in traffic was unprecedented but did not exceed the amount of time “runway” operators routinely plan for.

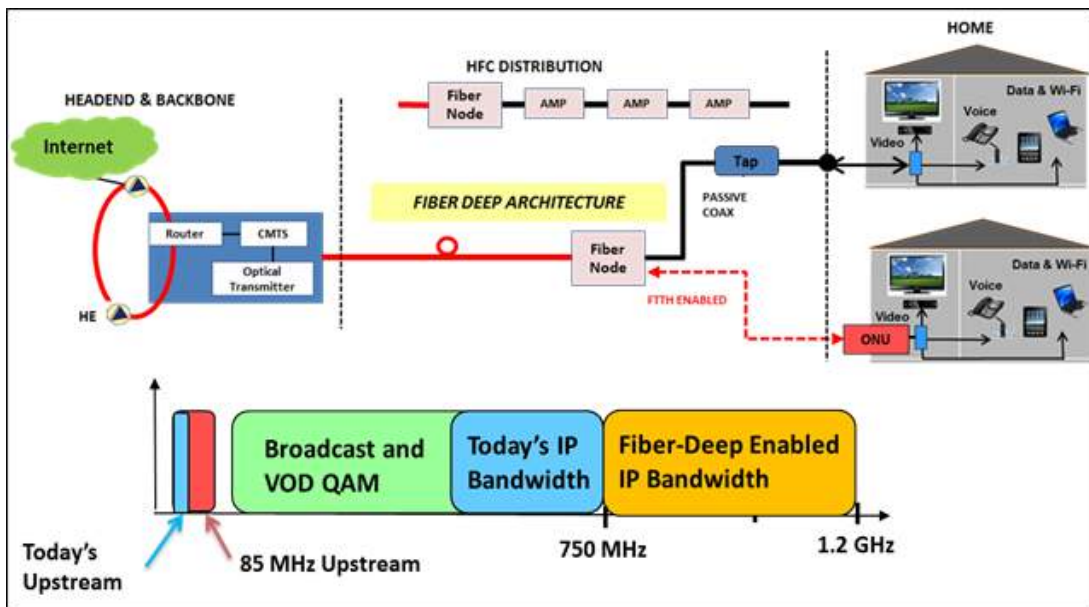
Despite pre-planning, the magnitude of the spike did create some of the aforementioned housefires where, for example, a portion of the network may have had an upgrade cycle a number of years ago, and at the time of the onset of the pandemic was not far from its next upgrade. In a “BAU” cycle there was more than sufficient capacity to make it to the scheduled upgrade – typically a node split. Capacity management teams carefully project “by when” an upgrade must take place or else risk a utilization/congestion alarm, by monitoring the network traffic at all sites and for all service groups, and building dashboards to make ongoing decisions, as shown in Table 1.

**Table 1 - Capacity Management Dashboards Guide Targeted Network Upgrade Investment [1]**

|        | Month 1            |                        | Month 2            |                        | Month 3            |                        | Legend - % of CMTS SGs |
|--------|--------------------|------------------------|--------------------|------------------------|--------------------|------------------------|------------------------|
|        | % >80% Utilization | Count >80% Utilization | % >80% Utilization | Count >80% Utilization | % >80% Utilization | Count >80% Utilization |                        |
| Site A | 0.02%              | 2                      | 0.04%              | 5                      | 0.17%              | 22                     | 0.00%                  |
| Site B | 0.05%              | 9                      | 0.13%              | 22                     | 0.49%              | 80                     | <.05%                  |
| Site C | 0.21%              | 33                     | 0.43%              | 67                     | 1.65%              | 251                    | <.1%                   |
| Site D | 0.09%              | 13                     | 0.06%              | 9                      | 0.33%              | 49                     | <.5%                   |
| Site E | 0.00%              | 0                      | 0.02%              | 2                      | 0.10%              | 11                     | <1%                    |
| Site F | 0.44%              | 20                     | 0.93%              | 42                     | 1.47%              | 66                     | >1%                    |


### 5.1.1.1 Rich Fiber Diet

Like all MSOs, Comcast continues to split nodes as traffic grows – it is a great advantage of the HFC architecture, that it is so incrementally scalable in a success-based way. However, several years ago, Comcast embarked on a network upgrade path that complements the typical node splits with a more aggressive approach that effectively acts as multiple node splits in a single upgrade. By design, this strategic plan will last for many years, in order to get ahead of the cycle of repeatedly going into the field in a matter of a couple of years. This is often referred to as a “Fiber Deep” or node plus zero (“N+0”) architecture. Millions of households passed have been built in this manner since introduction of this adjustment to the evolution plan was introduced. Much has been written about this architecture [5,6], which is shown in Figure 12.



**Figure 12 - N+0 Architecture Basics: No RF Amplifiers and Expanded DS and US BW**

A fundamental premise of Fiber Deep that led to its adoption is that it could be shown in certain areas to provide not just a simpler, higher performing, more consistent, and generally better last mile architecture of increased bandwidth, but it was also more cost effective over time than the series of node splits that would otherwise be needed to manage the traffic, as shown in Figure 13. It is not shown in Figure 13 but the “breakeven” time frame between those two options was about 7 years.



**Figure 13 - Fiber Deep Cost Effectiveness Compared to Continued Node Splitting Only [6]**

In summary, Fiber Deep is essentially the equivalent of executing multiple node splits at once, while expanding the spectrum in both directions. The architecture decreases the size of the node service group by typically 4-5x, while nearly tripling the US capacity and adding 60% more DS bandwidth.

Most notably, as the pandemic struck, utilization issues above thresholds marking them for augmentation were nearly non-existent in the footprint where this architecture was in place, owing to its proactively taken large bandwidth step. These networks will not require follow-up upgrades for many years, and thus were generally not under duress, despite the traffic spikes caused by COVID-19. From the CAGR and timeline perspectives, an architecture installed with a lifespan of 7-10 years holds up pretty well, even if a year, or 10-14% of it, is eliminated.

### 5.1.1.2 Widen the Lanes


The Fiber Deep architecture identified increases the available spectrum in both DS and US. The US expansion takes the available bandwidth from 37 MHz up to a limit of 80 MHz, also known as the Mid-Split architecture. It was defined in DOCSIS 3.0, with the upper limit selected in part to fall just below the FM radio band in the US, while preserving the important DS video out-of-band signals used by legacy QAM set-top boxes (STBs) widely deployed today. Per the earlier discussion, it is typically the US that drives network upgrade activity.

While Mid-Split is deployed in current N+0 networks at Comcast, it can also be deployed in N+X HFC networks, by installing different diplex filters in the RF amplifier cascade as well as the node. Because of the average per-user peak-busy-hour US is only in the hundreds of kbps range, US generally grows

slower than DS, and the new US spectrum is much cleaner, the impact on network lifespan of mid-split is very powerful.

Figure 14 shows the time runway generated by the three options – node split, node split plus upgrade to Mid-Split, and finally N+0. While N+0 offers the longest runway of the three, it is clear that an N+X migration tied to a node split is also a very effective way to extend HFC lifespan to nearly 7 years. Even with a COVID-19 traffic hit, it is easily absorbed while maintaining a multi-year runway.

Another benefit of N+X with spectrum migration is that it addresses the COVID-19 inspired objective to “add capacity quickly”, when compared to N+0. With the COVID-19 spike eliminating months of lifespan margin, N+X upgrades bring more US bandwidth to the network quickly to reset the lifespan timeline, and also support the growth of applications and new service speeds. The naturally slower pace of deep fiber construction will leave areas unaided by this architecture for a period of time, and with the “time” erased due to the pandemic, alternatives such as drop-in HFC upgrades that are both fast and effective make a sensible COVID-19 remediation step. Having a diverse strategy, not one-size-fits-all, adds important flexibility to deal effectively with adjustments to COVID-19.



**Figure 14 - Upstream Lifespan Expansion Options [1]**

Lastly, with speed to increased bandwidth support in mind, coupled with the desire to push fiber deeper into the network whenever possible, adjustments to Fiber Deep architectures and practices are worth considering. For example, adding fiber in the underground network is an inherently slow process. However, given the freedom to allow a strategically placed amplifier (e.g. to allow an N+1 network), then lessening fiber construction, increasing the node size, and decreasing cost all speed the pace of the network upgrade and deliver the added bandwidth to more households passed (HHP)/year, if the long-term capacity benchmarks can still be obtained.

### 5.1.1.3 Every Bit Counts


Both Fiber Deep and upgrades to new frequency splits are plant operations, and as such require planning and coordination. This is because they carry regional dependencies associated with construction processes, permits, municipal intervention, etc. In addition, the ability to access some portions of the network during a pandemic itself is limited. Fortunately, there is another important tool in the DOCSIS

toolbox that can address bandwidth directly and quickly and can provide a significant measure of relief. The tool is the DOCSIS 3.1 Profile Management Application, or PMA. It is closely related to the DOCSIS 3.1 feature known as Multiple Modulation Profiles, or MMPs. It fits the bill well as a quick relief valve to “recover” from COVID-19 traffic effects, while architectural plans of larger scope take longer to come together.

First, recall two of the major new changes between DOCSIS 3.0 and DOCSIS 3.1:

- 1) Single Carrier QAM (SC-QAM) was replaced with Orthogonal Frequency Division Multiplexing (OFDM) and Orthogonal Frequency Division Multiple Access (OFDMA). DOCSIS 3.1 recognized that it was the right time to break the mold and adopt more modern and mature OFDM. OFDM has the unique ability to optimize capacity across an uncertain channel in a way that SC-QAM cannot, at least not as easily. Instead of a 6 MHz chunk of spectrum to manage the signal over, the DOCSIS 3.1 flavor of OFDM uses granular subcarriers of 50 kHz or even 25 kHz in making up the composite signal. OFDM allows for narrow slices of spectrum to be optimized, and the overall channel to be used at its most efficient.
- 2) More QAM profiles and more efficient QAM modulation options were introduced, enabled by very powerful new Forward Error Correction (FEC). The base DOCSIS 3.0 of 256-QAM was increased to up to 4096-QAM, accounting for 50% more efficient use of bandwidth. All half-steps between 256-QAM and 4096-QAM were also enabled (512-QAM, 1024-QAM, 2048-QAM). QAM orders higher than 4096-QAM were added, but not as MUST requirements, and the efficiency bang for the dB-buck for these profiles is not very attractive.

Why is this so important after many successful years of 256-QAM? Figure 15 shows the recorded range of signal-to-noise ratio for millions of DOCSIS 3.0 cable modems, taken as part of the early work to develop the DOCSIS 3.1 standard. What it clearly shows is that, with the new FEC of DOCSIS 3.1 applied, many modems could achieve better than 256-QAM performance, and therefore many bits-per-second could be left on the table. How many bits are thrown away would vary depending on the particular modem.




**Figure 15 - Multiple Modulation Profile Potential of D3.1**

Because of the characteristics observed in Figure 15, multiple modulation profiles (MMPs) were defined in the DOCSIS 3.1 spec, allowing clusters of modems seeing similar fidelity across the band to be formed

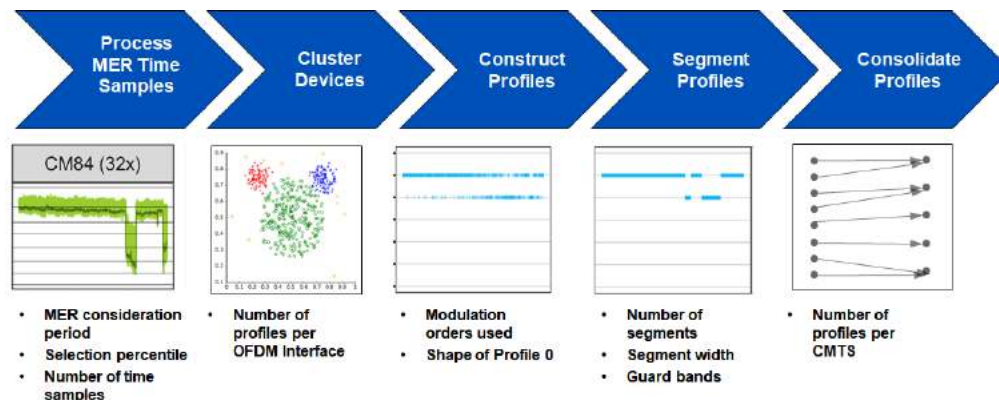


into groups that operate with common QAM modulation orders on their subcarriers. In the DS, each set of modems clustered by a common profile configuration would have their time on the wire per profile, as shown in Figure 16.



**Figure 16 - DOCSIS 3.1 Multiple Modulation Profiles vs Time and Freq [8]**


Figure 17 represents the workflow of fully capable, PMA-based capacity relief. It is coming to market in a very timely fashion to recover capacity taken away by COVID-19, and ensure that going forward the most efficient use of the channel is always available for the expected growth ahead, and furthermore executed as efficiently as possible with minimal human intervention. With MMP rather than a “flat” profile, for example, of all modems set at 1024-QAM, MMP fueled by PMA allows different modems to have a range of QAM profiles based on MER value. PMA brings an automated implementation, rather than thinking of MMP in a static sense or as a configuration set by the Headend Technician. PMA automates the process and allows for periodic checks and, if warranted, pushes updates to the modulation profiles without technician intervention. PMA will continually optimize algorithms as DOCSIS 3.1 continues to scale and big data metrics reveal more insights into temporal fidelity behavior on the coaxial plant.



**Figure 17 - Programmable Modulation Application Workflow [12]**

While the example above was based on the DS, the DOCSIS 3.1 US uses OFDMA, which has the same fundamental signal structure as OFDM. Like the DOCSIS 3.0 US, however, it must schedule and grant time-sliced access to the spectrum. Nearly all of the same logic of the DS applies with respect to optimizing the US channel in DOCSIS 3.1. However, most operators still have the majority of their network deployed with DOCSIS 3.0 modems, with DOCSIS 3.1 devices in the field supporting the higher DS speed tiers. Because of this larger percentage of D3.0 modems in the field, adding DOCSIS 3.1 in the US does not move the needle much to relieve capacity, since most devices cannot use it. Fortunately, with the US feeling more pressure from the COVID-19 traffic spike, and generally being more vulnerable, there are still DOCSIS 3.0 options available for capacity increases.

A common US configuration for operators is four bonded US 64-QAM carriers, each 6.4 MHz wide, which fills most of the high quality spectrum of a 5-42MHz US. However, there is additional spectrum, perhaps lesser in quality, that can be activated for more bps. Example use of the “5th and 6th” upstream in what had generally been empty spectrum is shown in Figure 18.



**Figure 18 - Exploiting More DOCSIS 3.0 Upstream in 5-42 MHz Systems**

Figure 18 suggests the presence of about 19% more useable spectrum, so in a 20% CAGR example would buy almost a year of time to work through a COVID-19-like spike. However, the spectrum available is likely to be impaired more than the existing 6.4MHz carriers, either due to filter roll-off and group delay variation at the high end, or low MER (noise aggregation) at the low end. So, the net capacity gain may be less than the spectrum gained, and the time bought by it shorter – but still meaningful for addressing a short-term crunch like COVID-19 while planning larger scale remediation steps.

Also, the DOCSIS 3.0 US offers several (four) settings for FEC overhead, where higher overhead means stronger FEC, and better protection for the packets on the DOCSIS 3.0 channels. DOCSIS 3.0 introduced 64-QAM about ten years ago, and, thinking in terms of in HFC network quality vs time, it was naturally set to maximum protection at launch, and rarely looked at since. However, since that time, node splits have shrunk service group sizes, and amplifier cascades have shortened. Initiatives that significantly enhance fidelity such as Fiber Deep and Distributed Access Architecture (DAA) are underway and in production at scale. As a result, some of those US channels, in particular the ones in the 25-40 MHz range, may no longer need maximum FEC protection. As such, there is a percentage of extra payload bits that can be gained in exchange for FEC overhead, depending on the quality of the highest MER US channels, from about 6% to, realistically, about 17% on a particular channel.

We will describe more about it in the following section, but modulation profile efficiency benefits considerably from the migration to DAA, which is being built in Fiber Deep markets and moving into N+X configurations also. The 10GbE digital fiber connectivity used in DAA significantly enhanced DS fidelity, by removing the analog optical links that typically had set the MER limit. It also enables higher-fidelity US signals to be present at the digital receiver. These higher MERs translate to more bandwidth efficient QAM signals.


In a situation like COVID-19, where a rapid response is essential, prioritizing software-based network enhancements, like PMA for capacity relief, can be executed and deployed generally more quickly than construction projects. They become exceptionally valuable tools for low-touch capacity augmentations, well suited to the kind of immediate response needs produced by the pandemic.

### **5.1.2 Be There for Me Always**

One of the key principles of the Fiber Deep architecture described is the simplification of the network. It reduces the number of active devices in the plant, removes the oldest actives in the plant, and provides a repeatable template for a last mile architecture that breeds consistency of process, equipment, RF level and fidelity performance, and aligns with expectations of future network applications and services.



Figure 19 shows a large sample of CMs taken across N+0 deployments. When compared to Figure 15, it is clear that the average MER has increased by 3-4 dB. Note that MER is slightly different than SNR, in that MER considers not just the noise degradation added to the signal, which SNR technically represents, but all of the other impairments as well. However, they have, unfortunately, come to be used interchangeably in the industry. Also, the bell-shaped curve seen previously is now favorably skewed towards higher MERs. In fact, the red arrow identifies the approximate average MER in the Figure 15 BAU N+X case when compared to the N+0 network. With the combination of significantly higher fidelity and the PMA tool, the network is well positioned to provide a higher quality, more consistent customer experience and effectively serve the increased volume of traffic the pandemic has wrought.




**Figure 19 - MER Distribution for a Large Sample of CMs on a N+0 System**

Indeed, trouble calls run consistently lower on the Fiber Deep parts of the network – ranging from 10-15% for the calendar year 2019, confirming the expectation that it is not only a superior network, it is also more cost effective to operate, and causes less disruption for busy WFH employees and online students.

Figure 20 shows a characteristic view of trouble call (TC) tracking in one Comcast Division that was very actively building N+0 in 2019. Note that the onset of COVID-19 during the initial “shock and awe” of shelter-in-place restrictions froze everything in place for a period of time (3/26/20) as companies and customers focused on getting a wide range of new priorities and changes in their day-to-day.

Not very long after – broadband being among the highest priorities – a return to normal TC interaction was taking place – but note this return to “normal” is with significantly increased COVID-19 traffic, now more essential traffic, and high awareness to network behavior among customers in the early days of the pandemic response.



**Figure 20 - Reduction of Trouble Calls in Fiber Deep Deployments**

In addition to prioritizing dense fiber access as an architecture, Comcast has pivoted to deploying Distributed Access Architecture (DAA) at scale, based on Remote PHY technology, using Remote PHY Device (RPD) Nodes from multiple established vendors. In conjunction with the transition to DAA, we launched a Virtual CMTS (vCMTS) core, leveraging 10 GbE connectivity from facilities to the RPD nodes in the field. The vCMTS is built as a containerized set of micro-CMTS cores, such that each RPD node has its own packet processing compute and is arrayed in a robust redundancy scheme with a very small blast radius, increasing availability and redundancy to guarantee the always-on WFH experience. The 10GbE connectivity sets the network up well to be last-mile access agnostic by implementing a range of last mile access technologies fed by simple Ethernet links deep into the plant.

The introduction of DAA and vCMTS delivered orders of magnitude more real time visibility into the network via streaming telemetry, and with it the ability to process advanced metrics and use machine learning in operations. The increased visibility is essential to deliver the proactivity critical to network availability, enabling tools and views to quickly assess and make determinations about anomalous network behavior, raise alarms, trigger self-healing and re-routing as necessary, and capture detailed logs to continue to learn about and optimize network robustness and flexibility in modern ways. Proactive visibility and availability are, of course, among those high priority areas of focus in support of the WFH and online education trends now and ahead.

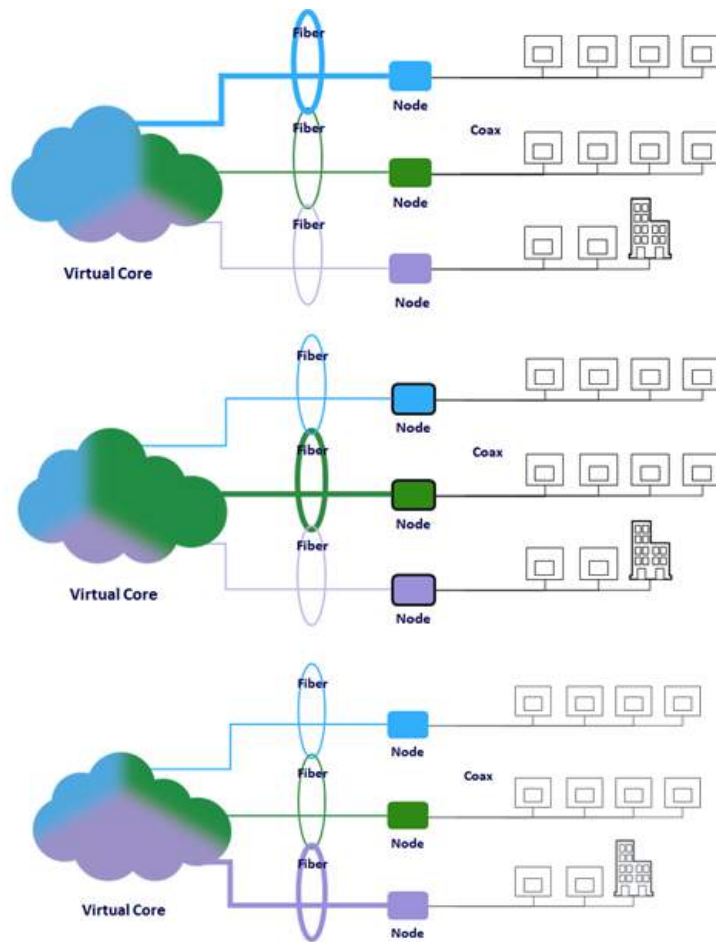
A sample dashboard summarizing the network status, with visibility on the order of seconds, is shown in Figure 21, encompassing information on CMs, RPDs, switches and routers. Virtualization in operator access networks leverages years of development and maturation in large scale data centers, driven by the likes of Amazon, Google, Facebook, etc. These providers, of course, have very stringent requirements for availability and reliability of their services, much like network operators do. While there are some important differences between operating a large data center that is handling “horizontal” data flows, compared to operating a network directly serving customers, the technological roots and the ecosystem scale created in establishing these data centers provides much of the foundation for virtualizing the access network edge.



**Figure 21 - Real-time Dashboard of Network and Component Availability and Uptime**

Another benefit of the virtualized core is its ability to throw packet processing compute in the direction where it is needed most, as shown in Figure 22. The pre-pandemic common example was daytime business services traffic in one part of the network geography, shifting to distributed residential web traffic during peak busy hour of the evening at home. In other areas of the footprint, such as universities and in apartment complexes tuned to a demographic, studying, gaming, chatting, or streaming late into the night is more the norm. The nimbleness of a virtual cloud could fall under the category “agile capacity” noted in the prior section. It landed here to separate this technology directional shift towards DAA and virtualization from the RF capacity focus of the prior section.

Enterprise services for office complexes are obviously not viewed as a segment that is expected to see a return to growth imminently. However, a network looking ahead should be considering the changes to business services and what might be their renewed directions of focus and investment. As implied in the example of Figure 22, business services traffic is a growing and increasingly important part of a cable operator’s services and revenue. Business for mid-size and large commercial enterprises often have specific, quantified, requirements for throughput, latency, jitter, packet loss, and availability. While the coaxial network is optimized for residential use, it does support a wide range of small-to-medium businesses (SMBs) with DOCSIS-based cable modems, typically with specific software loads that include certain features important for businesses, but with more limited Service Level Agreement (SLAs) than enterprise class services. And, while the coaxial system is undergoing development to improve availability and resiliency, as well as latency and jitter (to be discussed) – all key aspects of post-COVID applications – larger scale enterprise services are still better served by fiber-based Ethernet architectures.



**Figure 22 - The Virtual Cloud Efficiently Shifts Capacity Where it is Needed**


A further key principle behind the Fiber Deep architecture is that it is an investment in a platform that supports the growth in business services, many of which are fiber-based. With fiber reaching to less than 1000 feet of every coaxial connection, there will be increasing access to fiber for nearby business services. Furthermore, the migration to a DAA architecture means that this fiber infrastructure is becoming digital fiber based on 10 GbE. As the residential network evolution takes place, a more scalable, resilient, high visibility network is being put in place. On it, Layer 2 Ethernet services can share the wavelength space, the latter of which is multiplied substantially when going from traditional analog fiber optics, and 16 wavelength limitations, to the DWDM grid supporting up to 80 wavelengths. This 5x bandwidth multiple into these dense pockets has the network well-prepared for post-pandemic applications and traffic growth, such as commercial opportunities in industrial automation, M2M, and digitization applications tied to the COVID consequences on the manufacturing economy.

### 5.1.3 Be Snappy

Latency and jitter have historically applied primarily to voice services, before becoming dominated by the increasingly popular gaming segment. Gamers are no longer kids hiding in their parent’s basements. They are professionals who still love their games, and the games have become ever more elaborate and globally interactive. There have been many studies highlighting the sensitivity of games of certain types to network performance. Games generally do not generate high traffic volumes, but the customer experience can be significantly impacted if the packets are held up excessively between the gamer and the game server (or host). The bandwidth statement above does not necessarily apply to the category of cloud gaming or Augmented Reality/Virtual Reality (AR/VR) scenarios.

Of course, yesterday’s voice sensitivities to real-time attributes is today’s video conference. A major difference is that while there remains a sensitivity to latency and jitter, the bandwidth of these applications is much higher because they are video streams, and in a “Teams” meeting, for example, there can be a streaming video for each attendee using a camera. So, while capacity, bandwidth, and speed always dominate the discussion of access network evolution, latency is as much a factor if not more to the customer experience, and certainly becomes magnified in the pandemic-driven video conference world. Speed represents the amount of data that can be sent across a connection over a fixed period of time, whereas latency is the time it takes for a packet to cross the network and get a response. It is the “ping” time.

Seeing latency as an emerging priority for real time services such as gaming, video conferencing, IoT M2M, etc., working to address the components of the latency budget attributed to the DOCSIS network is an industry priority. A program was initiated at CableLabs to standardize the approach across the technology ecosystem. Obviously, the added impetus given by WFH has only increased the interest in developing Low Latency DOCSIS (LLD). In the DOCSIS 3.1 network, the dominant causes of latency are queuing and media acquisition [13], as shown in Figure 23.



**Figure 23 - Sources of Latency in DOCSIS 3.1 Networks [13]**

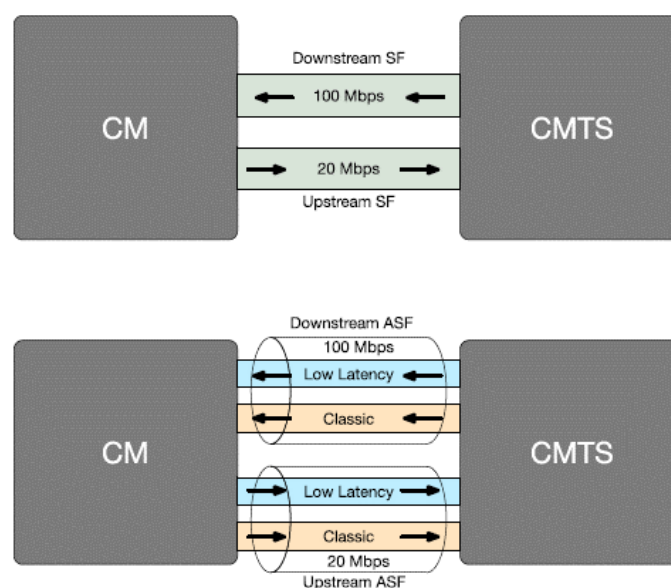
The lowest hanging fruit is in the area of queuing. Once controlled, there is a potential to drive latency into single-digit milliseconds in the “media acquisition” area. The term “media acquisition” in this case refers to the request-grant mechanism by which a CMTS makes determinations about access to the upstream channel time slices and provides that information to the modems. Low millisecond values are instantaneous in “human” time and open the door to future IoT and machine-to-machine communications



where millisecond and sub-millisecond delays can have meaningful impact, for example, in stock trading or industrial machinery.

The potential for major improvements by managing queueing more effectively is clear from Figure 23. These queueing drivers are primarily the nature of the Transmission Control Protocol (TCP), the Layer 4 mechanism used throughout the Internet. High bandwidth applications ramp to a rate faster than the network bottleneck can support, and rely on TCP congestion control, which inform the sender to back off when buffers holding the traffic bursts fill, at which point it begins to drop packets to manage the queue. Because of this, latency and packet loss are inherent in TCP links. LLD aims to improve upon this behavior.

LLD does not create a fast path and a slow path or prioritize traffic one way or another. The approach is to identify traffic characteristics in applications that contribute to building of the queues which create latency, and steer traffic to a second queue if it is a so-called Non-Queue-building (NQB) traffic. This “dual queue” approach, shown in Figure 24, shares the bandwidth allocation to the user amongst all of their active applications, and also means less packets dropped out of the queue-building buffer. It is easily enabled using DOCSIS Service Flow (SF) definitions, available today, that shape the aggregate traffic, and the individual flows defined by the dual queues have shaping turned off, avoiding prioritization. This approach improves the customer experience for applications that are latency sensitive without affecting the experience of the other services.




**Figure 24 - Traditional and LLD Service Flow Architecture [13]**

In addition, statistics are generated on the queueing delay by the CM and CMTS for use by MSOs, enabling them to fully understand the contribution of the DOCSIS network and help itemize their end-to-end latency budget. The customer experience can be closely correlated to these latency metrics. We already have tools in place that observe network latency further northbound in the network, and use this data, filtered and correlated across devices and bonding group assignments, to make proactive determinations to drive optimizations and customer QoE.

CableLabs is engaged with the IETF and the broader industry to identify the proper queue-using classifiers within the fields of IPv4 and IPv6 traffic. Most of today’s traffic volume is queue-building

(QB), in particular streaming video. These services transmit packets faster than the network bottleneck can support, relying on the TCP algorithms to buffer and alert the source to adjust accordingly. By contrast, NQB traffic sources generally do not send data faster than the network can directly support. Skype (and similar) and online games generally fall into this category, as does basic web browsing. Despite their inherent “good behavior” relative to utilization of network resources, their packets are unnecessarily mingled into the queue created because of the QB applications today.

The LLD effort includes additional options to further reduce latency, in particular to speed up the DOCSIS request and grant cycle of scheduling packet transmission US. These include shortening the window between request options and proactively granting access based on traffic trends and utilization, which can be learned over time, such that inefficiency is minimized. Figure 25 shows the potential for reducing the DOCSIS network latency as these features become available over time.



**Figure 25 - Industry LLD Efforts are Promising for Supporting Real Time Services Such as the Recent Explosion of Video Conferencing Due to WFH [13]**

As with other access network initiatives, existing technology development underway is aligned very well with addressing the impacts of the pandemic. The immediate WFH thrust and online schooling is certainly a reason MSOs have cause to consider adjusting priorities to move LLD higher on the list of To-Do's. These performance improvements also align well to broadband requirements that are expected to be part of industrial digitization.

## 6.0 Conclusion

One way to look at the COVID-19 situation as it applies to the network is to acknowledge two things that are somewhat in opposition to one another:

- 1) We have never before seen a traffic dynamic like this – the magnitude of the increase plus the abbreviated period over which it occurred – in our broadband cable lifetime
- 2) Simultaneously, there is little new that needs to be learned to figure out how to handle it.

We identified four areas of priority from which to re-calibrate access evolution plans, given the observed trends and anticipated permanent impacts on the network attributable to COVID-19:

*Enable Agile Coaxial Capacity* – Take advantage of existing DOCSIS 3.0 platforms to add new channels, optimize DOCSIS 3.0 US profiles, and DOCSIS 3.1 PMA to add capacity to recover the “time” lost by the COVID-19 traffic wave. These techniques complement more time-consuming operator network upgrade plans, such as changing the frequency split in the network by swapping actives, splitting nodes, and pulling fiber deeper.

*Hardening and Resiliency* – Synergistic with pulling fiber deeper, adding more nodes, and removing the oldest actives in the plant – to reduce trouble calls and increase availability – is a transition to DAA and virtualization of the CMTS core. Granular streaming telemetry keeps an eye on the network more closely and in real time than ever before, data analytics trigger proactive alerts and notifications, and the micro-CMTS architecture further maximizes availability and adds resiliency.

*Responsivity* – Drive down network-induced latency and jitter to support the increase in real-time applications, such as video conferencing and online classrooms, ensure the increased in gaming traffic does not impact the gamer experience, and to support new opportunities that arise in business services that have real-time and near-real time broadband implications.

*Enabling of Enterprise Services* – Change often leads to opportunity, and major change, on the COVID-size scale, will likely lead to a range of new industrial opportunities aimed and helping companies be prepared to “weather the storm” in the future. What may at first be a softness in business services opportunities is likely to take a favorable turn in the future, with perhaps a change in the types of business markets served. The fiber-rich, Ethernet-based DAA system has the flexibility to be well-suited to deliver a complete range of last mile solutions and capabilities that may be uncertain today.

If there was ever a scenario to be caught by surprise and in a desperate scramble to recover, this would have been that time. While there was indeed some scrambling, the broadband cable network gets an “A+” for its performance in the face of this unprecedented traffic wave. The reasons for this are many, and boil down to a few key points

- Planning several years ahead is the nature of the job for access network architects and capacity managers, and it paid off in this crisis
- Very familiar muscles can be exercised efficiently to deliver additional capacity, built upon many years of experience honing these processes
- Additional network technology was already in the works that was aimed at addressing the needs observed during the pandemic and what is expected to be ongoing beyond it



- Efficient-to-deploy, SW-based new technologies, such as PMA and LLD, are emerging at a very timely moment, because the benefits they bring had already been deemed important for future applications that happened to arrive at scale a bit earlier
- For specific applications of increased priority, the scale and penetration pace of them changed quickly, but there are not additional technology barrier dependencies
- Cable operators have been highly focused on the upstream very recently, because bandwidth from the home is close to filling the availability capacity
- Cable operators have been highly focused on network monitoring and availability as the cable product offering has shifted towards always-on Internet access as a top priority

Let's hope this is the last pandemic any of us witness in our lifetimes. However, having passed the network stress test of a generation, we can all be comforted by the recognition that the power of the broadband network is there to ensure that some semblance of "life goes on" is still achievable. And undoubtedly, there will be lessons learned that can make us even better prepared should such a situation strike again.

## Abbreviations

|        |                                               |
|--------|-----------------------------------------------|
| AR/VR  | Augmented Reality/Virtual Reality             |
| BAU    | Business as Usual                             |
| CAGR   | Compounded Annual Growth Rate                 |
| CM     | Cable Modem                                   |
| CMGR   | Compounded Monthly Growth Rate                |
| CMTS   | Cable Modem Termination System                |
| DAA    | Distributed Access Architecture               |
| DS     | Downstream                                    |
| ISBE   | International Society of Broadband Experts    |
| SCTE   | Society of Cable Telecommunications Engineers |
| FEC    | Forward Error Correction                      |
| HHP    | Households Passed                             |
| HSD    | High Speed Data                               |
| IETF   | Internet Engineering Task Force               |
| IoT    | Internet of Things                            |
| LLD    | Low Latency DOCSIS                            |
| MER    | Modulation Error Ratio                        |
| M2M    | Machine to Machine                            |
| MMP    | Multiple Modulation Profiles                  |
| MSO    | Multiple Systems Operator                     |
| N+0    | Node Plus Zero Amplifiers                     |
| N+X    | Node Plus “X” Amplifiers                      |
| NQB    | Non-Queue-Building                            |
| PBH    | Peak Busy Hour                                |
| PMA    | Profile Management Application                |
| QB     | Queue Building                                |
| OEMs   | Original Equipment Manufacturer               |
| OFDM   | Orthogonal Frequency Division Multiplexing    |
| OFDMA  | Orthogonal Frequency Division Multiple Access |
| QoE    | Quality of Experience                         |
| RPD    | Remote Phy Device                             |
| SC-QAM | Single Carrier QAM                            |
| SF     | Service Flow                                  |
| SLA    | Service Level Agreement                       |
| SMB    | Small-to-medium business                      |
| TC     | Trouble Call                                  |
| TCP    | Transmission Control Protocol                 |
| TPD    | Traffic Doubling Period                       |
| US     | Upstream                                      |
| vCMTS  | Virtual CMTS                                  |
| VPN    | Virtual Private Network                       |
| WFH    | Work from Home                                |

# Bibliography & References

- [1] Barker, Bruce E, and Claude Bou Abboud, Erik Neeld, “Access Capacity Planning: Staying Well Ahead of Customer Demand Helped Ensure Stability During COVID-19,” SCTE Cable-Tec Expo, Oct 13-16.
- [2] Bienkowski, Tom, “Pandemic Life Online: What is Everybody Doing,” [www.netscout.com](http://www.netscout.com)
- [3] Bienkowski, Tom, “COVID-19 Network Traffic Patterns: A Worldwide Perspective from Our Customers,” [www.netscout.com](http://www.netscout.com)
- [4] Branscombe, Mary, “The Network Impact of the Global COVID-19 Pandemic,” <https://thenewstack.io>
- [5] Howald, Dr. Robert L, Comcast, “The Fiber Frontier,” INTX Spring Technical Forum, May 16-18, 2016, Boston, MA.
- [6] Howald, Dr. Robert L, Comcast, “Aboard the Technology Wave: Surf Report,” SCTE Cable-Tec Expo, Sept 26-29, 2016 Philadelphia, PA.
- [7] Howald, Robert L, and Dr. Sebnem Ozer, Robert Thompson, Saif Rahman, Dr. Richard Prodan, Jorge Salinger, “What is 10G – The Technology Foundation,” SCTE Cable-Tec Expo, Sept 30-Oct 3, 2019, New Orleans, LA.
- [8] Howald, Robert L, , “Network Preparation: Maximizing Capacity ROI,” SCTE Cable-Tec Expo, Oct 21-24, 2013, Atlanta, GA.
- [9] Koeze, Ella and Nathaniel Popper, “The Virus Changed the Way We Internet,” [www.nytimes.com](http://www.nytimes.com)
- [10] Liberatore, Stacy, “Coronavirus has changed the way American use the Internet,” [Dailymail.com.uk](http://Dailymail.com.uk)
- [11] Moritz, Scott, “Empty Offices, Full Homes: COVID-19 Might Strain the Internet,” [www.bloomberg.com](http://www.bloomberg.com)
- [12] Rice, Daniel, “DOCSIS 3.1 Operational Hardening,” CableLabs Summer Conference, August 4, 2019, Keystone, CO.
- [13] White, Greg, and Karthik Sundaresan, Bob Briscoe, “Low Latency DOCSIS: Technology Overview,” <https://cablela.bs/low-latency-docsis-technology-overview-february-2019>.
- [14] COVID-19: How Cable’s Internet Networks Are Performing: Metrics, Trends & Observations, <https://www.ncta.com/COVIDdashboard>
- [15] (<https://www.cdc.gov/coronavirus/2019-ncov/cases-updates/cases-in-us.html>)

# Acknowledgments

The author would like to thank Brian O'Neill of CommScope for some enlightening discussion and insights regarding the view on COVID-19 impacts from the OEM perspective.

The author would also like to thank his indispensable administration professional, Karen Murray, for taking the lead on researching publications regarding the pandemic's effect on networks.

# Wi-Fi Passwords

## The Evolving Battle Between Usability and Security

A Technical Paper prepared for SCTE•ISBE by

**Craig Pratt**

Lead Software Engineer, Security Technologies  
CableLabs  
858 Coal Creek Circle, Louisville, CO 80027  
+1 303.661.3408  
c.pratt@cablelabs.com

**Darshak Thakore**

Principal Software Architect, Security Technologies  
CableLabs  
858 Coal Creek Circle, Louisville, CO 80027  
d.thakore@cablelabs.com

**Jacob Gladish**

Director of Software Development  
Comcast  
1701 JFK Boulevard, Philadelphia, PA 19103  
jacob\_gladish@comcast.com

# Table of Contents

| Title                                                                        | Page Number |
|------------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                         | 3           |
| 1.1. One Password to Rule Them All.....                                      | 3           |
| 1.2. The Evolution of Wi-Fi Credentials.....                                 | 3           |
| 1.3. Goals for the Next Generation of Wi-Fi Authentication and Security..... | 4           |
| 2. Wifi Device Authentication and Key Establishment .....                    | 4           |
| 2.1. WPA2 (802.11i-2004).....                                                | 4           |
| 2.1.1. WPA2 Personal.....                                                    | 5           |
| 2.1.2. WPA2/WPA3 Enterprise .....                                            | 7           |
| 2.1.3. Other WPA2 PMK Establishment Methods .....                            | 9           |
| 2.2. WPA3-Personal.....                                                      | 11          |
| 3. Provisioning Station Credentials using DPP .....                          | 12          |
| 3.1. DPP Provisioning .....                                                  | 12          |
| 3.2. DPP/EasyConnect Connectors.....                                         | 15          |
| 3.3. Applications of DPP Connectors .....                                    | 17          |
| 4. Wireless Network Segmentation .....                                       | 18          |
| 4.1. Network Segmentation using DPP Connectors .....                         | 20          |
| 5. Conclusion.....                                                           | 20          |
| Abbreviations .....                                                          | 21          |
| Bibliography & References.....                                               | 21          |

## List of Figures

| Title                                                               | Page Number |
|---------------------------------------------------------------------|-------------|
| Figure 1: The Wi-Fi/EAPOL 4-Way Handshake.....                      | 5           |
| Figure 2: PTK establishment in WPA-Personal .....                   | 6           |
| Figure 3: Key Usage in WPA2-Personal.....                           | 6           |
| Figure 4: WPA2-Enterprise Authentication .....                      | 7           |
| Figure 5: PTK Establishment in WPA2-Enterprise.....                 | 8           |
| Figure 6: Key Usage in WPA2-Enterprise.....                         | 8           |
| Figure 7: PTK Establishment Using MAC-associated Password/PMK.....  | 9           |
| Figure 8: PMK Distribution using MAC-associated Password/PMK .....  | 10          |
| Figure 9: WPA3 Personal SAE Exchange .....                          | 11          |
| Figure 10: DPP Direct Provisioning of Connector Using QR Code ..... | 12          |
| Figure 11: DPP Direct Provisioning of PSK Using QR Code.....        | 13          |
| Figure 12: DPP AP-based Provisioning of PSK using QR code.....      | 14          |
| Figure 13: DPP Provisioning of PSKs.....                            | 15          |
| Figure 14: DPP Connector Example.....                               | 16          |
| Figure 15: Encoded DPP Connector Example.....                       | 16          |
| Figure 16: DPP Connector-based Authentication .....                 | 17          |
| Figure 17: Example Network Segmentation .....                       | 19          |

# 1. Introduction

By all measures, Wi-Fi is a phenomenal success. As broadband availability and capacity increased, the speed of Wi-Fi also increased. As the number and nature of devices supported by Wi-Fi expanded, the range and capabilities also expanded. And when Wi-Fi entered the corporate/enterprise workspace, accommodations were added to allow for Wi-Fi credentials to be provided via integration with enterprise-level user authentication systems.

## 1.1. One Password to Rule Them All...

For most Wi-Fi networks, credential provisioning comes down to the question: “What’s the Wi-Fi password?”, followed by the tedious task of verifying the spelling and punctuation. And while there have been various mechanisms created to share Wi-Fi passwords across devices, the vast majority of Wi-Fi networks utilize a single shared password.

The implications of having a single credential for a Wi-Fi network have been known for years. With the first version of Wi-Fi security (Wired Equivalent Privacy or WEP), the fact that there was a single persistent key for the network (among other issues) led to a serious security vulnerability that would allow the key to be derived by simply observing a sufficient amount of network traffic. WEP password cracking was addressed in WFA Wi-Fi Protected Access 2 (WPA2) [3]. But even in WPA2, one authorized Station can observe the traffic of any other Station by observing its authentication exchange. And one issue is inescapable: once a password is shared, it cannot be unshared. The only remedy today for revoking a station’s credentials on a single-password networks is to change the password and reprovisioning all devices that need access with the new password.

While reprovisioning all Stations on a Wi-Fi network might have been a practical (if not tedious) task in the past - with perhaps a half-dozen interactive devices on the network (laptops, smart phones, etc.) - today it is a virtual impossibility. IoT devices in particular have proliferated enormously in recent years and are notorious for having inconsistent means of provisioning Wi-Fi credentials. And their interfaces and documentation workflows are geared for initial “quick start” setup, not reprovisioning. Many devices even require a complete factory reset to be provisioned – leading to a cascade of issues with device reconfiguration, updating apps, and reassociating devices with their various cloud services.

## 1.2. The Evolution of Wi-Fi Credentials

There have been a variety of solutions proposed and implemented to help mitigate the single password issue: Multi-zone networks (e.g. “guest” networks) allow for one password to be used for resident users/devices and another for non-resident users/devices. However, the way guest networks are implemented requires multiple Wi-Fi networks (SSIDs) to be configured on a wireless Access Point (AP) – which either dedicates separate Wi-Fi channels to the guest network or increases the amount of overhead on shared channels. And as soon as the password for the resident zone is shared with a device or user, the zone can be compromised.

WPA2/3 Enterprise was designed to support the integration of enterprise environments utilizing centralized user provisioning systems with Wi-Fi provisioning. These user-level authentication systems were tailored for interactive enterprise devices (e.g. laptops and mobile devices). This allowed the use of enterprise management systems to manage Wi-Fi access. The same aspects that make WPA2/3 Enterprise well suited to enterprise environments, however, make it difficult to use for home/SOHO environments. For instance, while it can be natural to provision interactive devices such as laptops and mobile devices with user credentials, IoT devices are not as easily associated with a single user – and almost universally

lack support for WPA2/3 Enterprise. And the infrastructure required for WPA2/3 Enterprise support is too complex for the average user to configure and manage.

MAC address filtering solutions attempt to provide per-device in a different way: While a user/device may gain access to the network using a shared password credential, other components of the Wi-Fi Access Point can filter traffic based on the MAC address of individual device(s). Revoking “access” to a device can be done at-will and based on various criteria – such as time of day, the internet host(s) being accessed, or metering criteria (e.g. time on-line or data usage). These solutions are currently employed by many network providers and are the bane of teenagers across the Internet. But they are also easily defeated by credential sharing and MAC spoofing – and they don’t prevent the observation of network traffic using the shared credential.

### **1.3. Goals for the Next Generation of Wi-Fi Authentication and Security**

Ultimately what is needed for truly manageable Wi-Fi networks is a mechanism that allows each device to have a unique, verifiable identity and allows for the provisioning of credentials which are cryptographically strong, unique and revocable. And to help ensure adoption and make everyone’s life easier, provisioning of devices should be as easy – if not easier – than the single-password method. Once individual identities and credentials can be associated with Wi-Fi devices, a new world of possibilities opens up for Wi-Fi network management and device/network security.

One solution that has been in development in the Wi-Fi Alliance is the Device Provisioning Protocol (DPP) specification – also referred to as “WFA Easy Connect”. The DPP specification enables the secure transfer of device credentials and metadata on both home and enterprise networks using a variety of provisioning mediums and mechanisms without complex infrastructure and without entering password/pre-shared keys (PSKs). In this paper we describe how DPP/Easy Connect works, the context and justification for it, and how it can be built upon to provide more safe and secure Wi-Fi networks while also enabling the advanced networking features we need in the future.

## **2. Wifi Device Authentication and Key Establishment**

### **2.1. WPA2 (802.11i-2004)**

The first generation of Wi-Fi security (WEP) required all the wireless-enable devices (“Stations”) to share a common 40- or 104-bit encryption key. The Access Point and Stations were all manually configured with the same 10- or 26-digit hex-coded key. This key was used directly to encrypt every payload packet sent across the network. Fatal security issues were discovered in WEP that led to it being deprecated – and replaced by WPA. But security issues aside, there was also a need for more flexible per-Station key establishment methods than either WEP and WPA could provide in order for Wi-Fi to be adopted in more environments.

To allow for more flexibility in how keys are associated with Stations, the WPA2 specification (802.11i-2004) introduced “pairwise” and “group” keys. Specifically, WPA2 defines the following keys:

- **The Pairwise Transient Key (PTK):** This key is used to encrypt all unicast (non-broadcast/multicast) traffic between one Station and the Access Point. The PTK is established using the “4-Way Handshake”
- **The Group Transient Key (GTK):** This key is used to encrypt all broadcast and multicast traffic on the network attached to the AP. The GTK is established and distributed using the “Group Key Handshake”



- **The Pairwise Master Key (PMK):** This key is used, along with other parameters, to establish the PTK. The method for establishing the PMK varies according to the authentication method.

All Wi-Fi key establishment methods use the 4-way handshake to establish the PTK and GTK for a Station. What varies between the methods is the establishment of the PMK, which will be described in later sections.

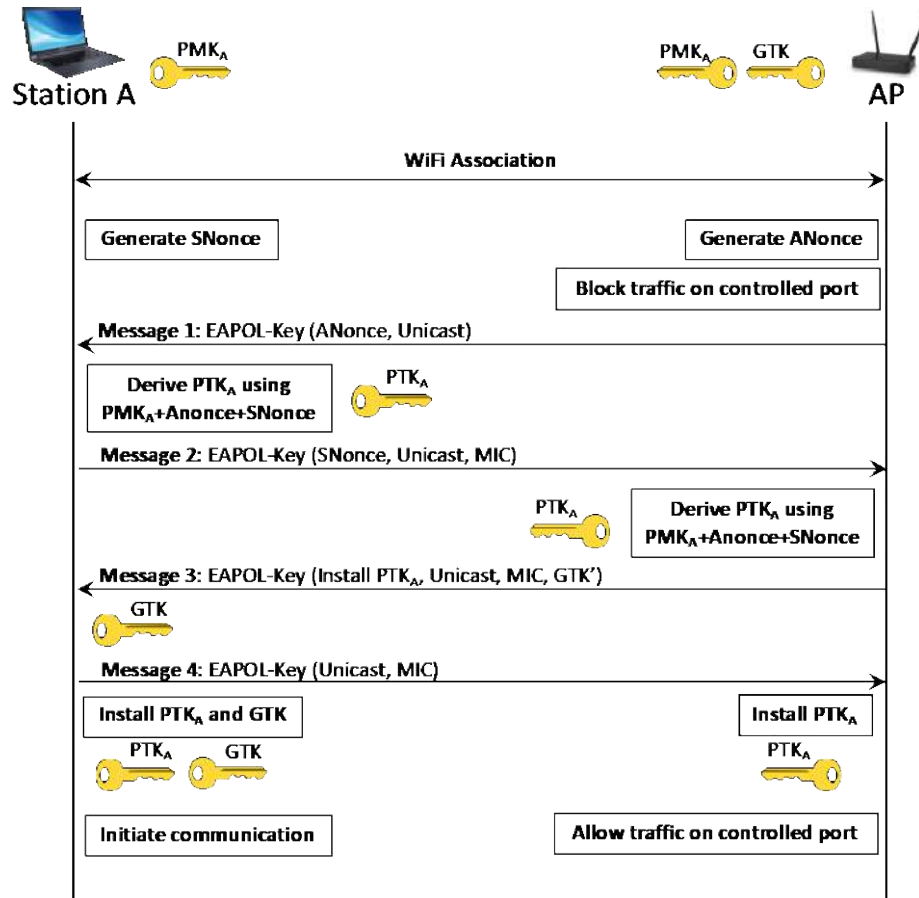
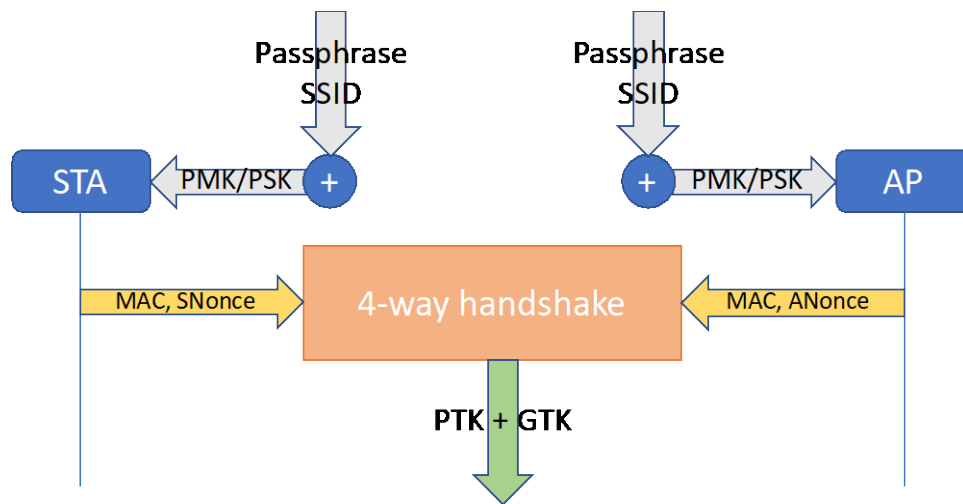


Figure 1: The Wi-Fi/EAPOL 4-Way Handshake

### 2.1.1. WPA2 Personal

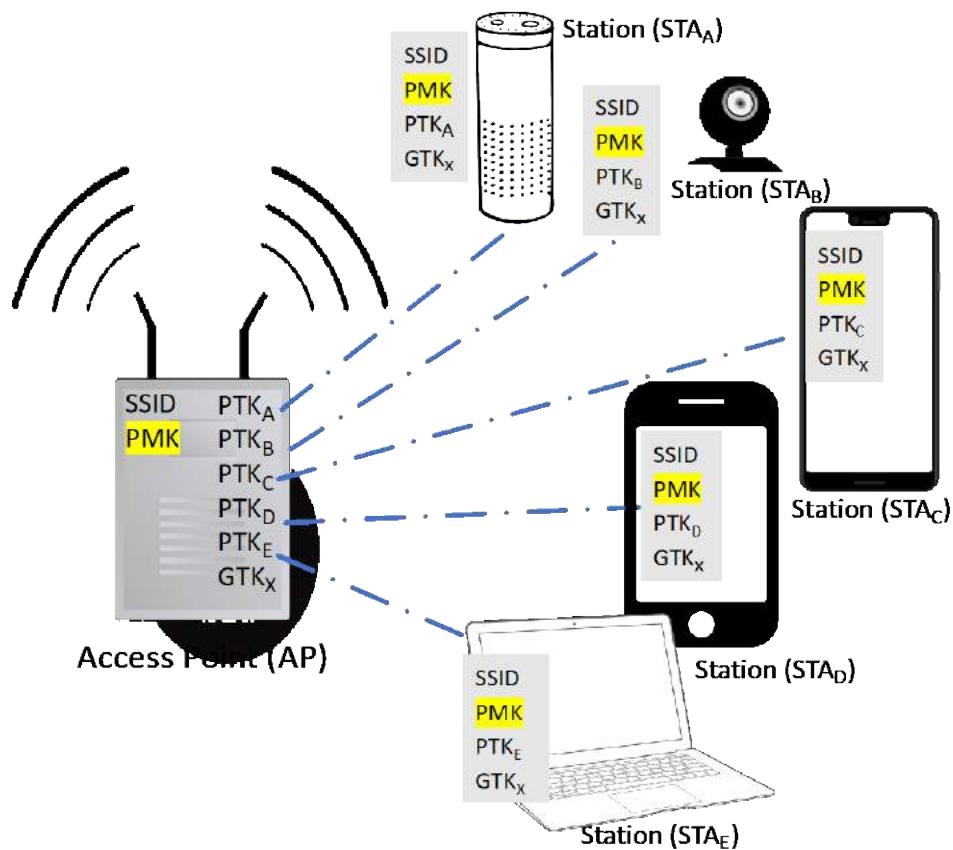
Residential, small/home office (SOHO) networks often use the WPA-Personal profile – which utilizes a shared password to determine the PMK. All that is required to establish a WPA-Personal Wi-Fi network is to configure the AP with an SSID and password and configure all the Stations with the same password.

The 4-Way Handshake allows the Station and the AP to establish that they both have the same password and allows for the establishment of a per-Station PTK.



**Figure 2: PTK establishment in WPA-Personal**

What's notable here is that – while each Station and the AP will have the same PMK (by virtue of using the same password) – the 4-Way Handshake ensures that each Station has a *unique* PTK (by virtue of the fact that the Snonce and Anonce are randomly generated). The GTK, on the other hand, is the same for all entities on the network so that shared traffic (such as broadcast and multicast) can be decrypted by all parties on the network.



**Figure 3: Key Usage in WPA2-Personal**

Note that every time a Station reauthenticates, a new PTK is generated. And any time a Station leaves the network, a new GTK is generated by the AP and distributed to the currently authentication Stations to ensure that Stations which have been removed from the network cannot decrypt group traffic using a previously derived GTK.

While WPA2 Personal is simple to setup – and simple to authenticate – this simplicity comes with a price: any Station(s) with the active password can authenticate. And the only way to prevent an unwanted Station from authenticating is to manually change the password on the AP and all the Stations – which is highly disruptive, especially with IoT devices. Additionally, any Station with knowledge of the Password/PMK can derive the encryption key of any other Station by observing the authentication exchange of that Station and the AP.

### 2.1.2. WPA2/WPA3 Enterprise

As mentioned previously, shared passwords are not an appropriate authentication solution for environments where user/Station access needs to be tightly controlled. Specifically:

1. In enterprise environments, there is typically already an authentication system in place – and corporations wishing to deploy Wi-Fi want to use their existing systems to authenticate Stations using existing user credentials.
2. The level of access to an enterprise network is often user-dependent. A shared password does not allow for differentiated or revocable Wi-Fi service. Some form of user identity is required. E.g. The user ID can be used to establish which VLAN (Virtual Local Area Network) and other resources the Station will have access to.
3. Revoking access to certain users/Stations should not require disrupting the Wi-Fi network and/or reconfiguring other Stations – as is the case with WPA2-Personal.

WPA2-Enterprise addresses the limitations of WPA2-Personal by providing a means of user-level authentication prior to the 4-Way Handshake using EAP-based (Extensible Authentication Protocol) authentication with a AAA server implementing the RADIUS/Diameter protocol. WPA3-Enterprise utilizes the same authentication methods and protocol as WPA2-Enterprise while updating the security profile of the various authentication methods and management messages. For example, WPA3-Enterprise requires Wi-Fi management frames to be protected and deprecates the use of cipher suites that are now considered insecure.


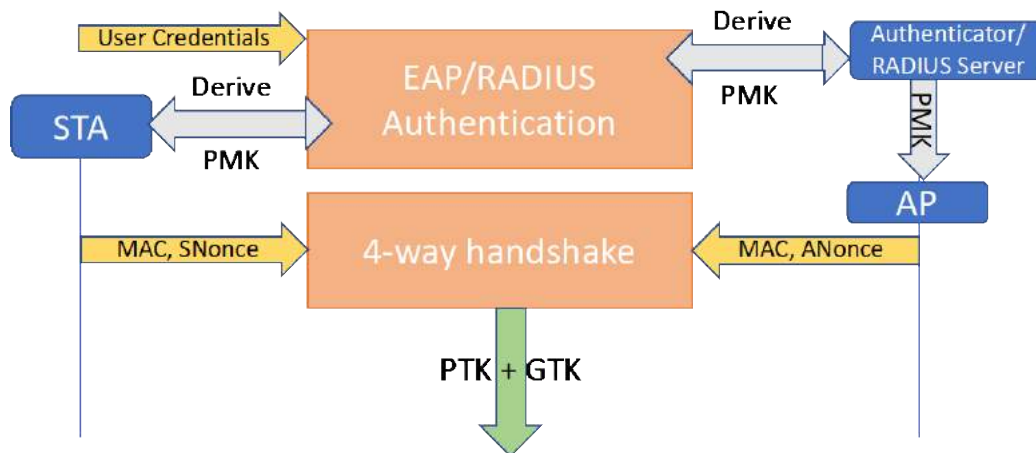
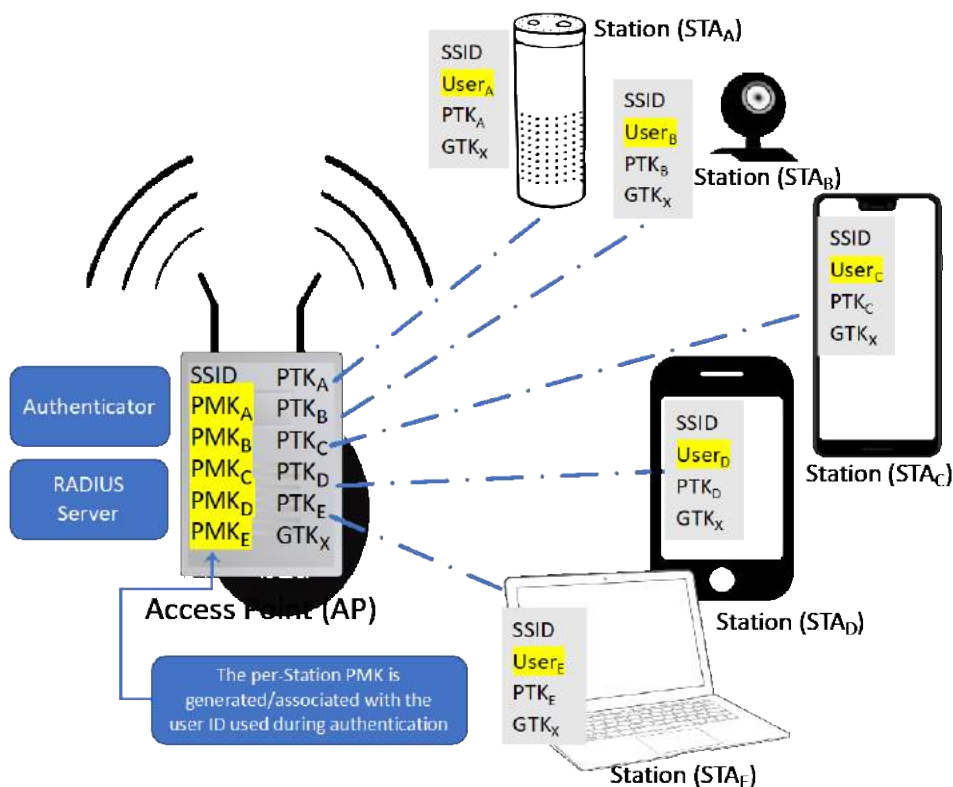


Figure 4: WPA2-Enterprise Authentication

Once the Station (“Supplicant”) and Authentication Server agree on an EAP method, they can exchange an arbitrary set of EAP messages to perform authentication. If authentication is successful, the Authentication Server and Station derive a shared key and the Authentication Server provides the derived key to the Authenticator. The Authenticator uses the EAP-derived PMK to drive the 4-way handshake with the Station.



**Figure 5: PTK Establishment in WPA2-Enterprise**



**Figure 6: Key Usage in WPA2-Enterprise**

WPA2-Enterprise addresses the single-password shortcomings of WPA2-Personal – and can satisfy the common enterprise requirements listed above when properly deployed and configured. But it does so at

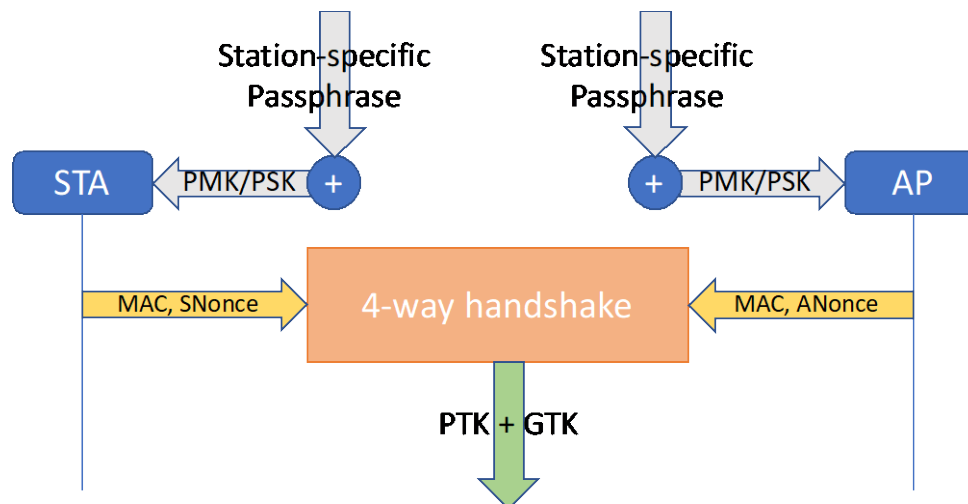
the cost of additional infrastructure, user/credential provisioning, and additional configuration. In particular, WPA2-Enterprise solutions are often tailored for user-based authentication using enterprise credentials. While this is well-suited to enterprise environments with interactive clients and a centralized user credential database, it's less suited for IoT and other dedicated-use devices.

User or device-level authentication can alternatively be performed in WPA2/3 Enterprise by using X.509 certificates. This form of authentication requires installing certificates and keys into the devices. However WPA2/3 Enterprise by itself does not provide a standardized mechanism to provision certificates/credentials and associated CA certificates into device trust stores.

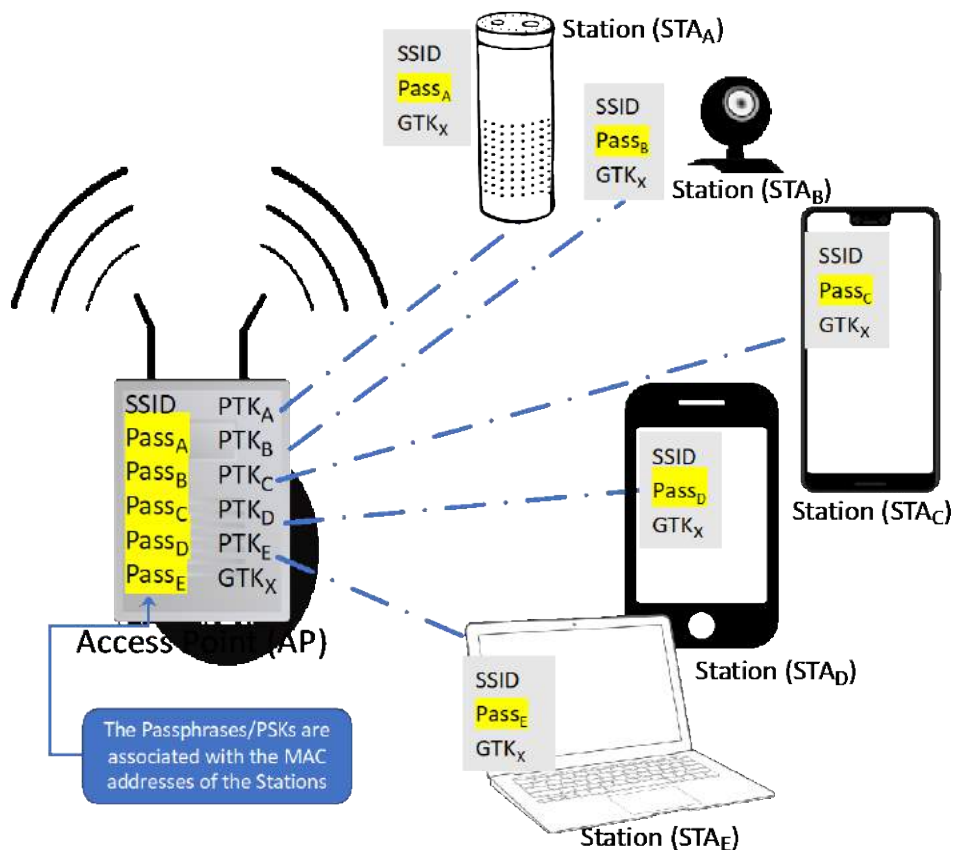
### 2.1.3. Other WPA2 PMK Establishment Methods

As noted in section 2.1.1, the simplicity of using a shared password in WPA2-Personal is also its main deficiency. And while WPA2-Enterprise enables per-device credentials, it introduces other issues. But as might be clear by now, any method that provides a per-station PMK can enable per-Station credentials, Station-specific policy, and Station revocability.

The key to deriving a unique per-station PMK is to determine an identity for a Station. The simplest form of identification for a Station is the MAC address of its Wi-Fi interface. Each Station can be configured with a unique WPA2-Personal password and the AP can associate a password with the MAC of the station. When the Station attempts authentication, the AP can look up the password associated with the MAC address of the authentication frames and initiate the WPA2 4-Way Handshake. The same MAC association can be used to establish per-Station policy (e.g. VLAN association, access policies, etc.).



**Figure 7: PTK Establishment Using MAC-associated Password/PMK**



**Figure 8: PMK Distribution using MAC-associated Password/PMK**

This MAC association method is supported by the reference Wi-Fi AP implementation *HostAPD* as well as a variety of commercial products/offerings.

There are some issues with this method however:

- The first time a Station connects with a password, the MAC isn't known by the AP. If the password/PSK is a single-station credential, the AP will learn the MAC address of the first Station that connects with the Password/PSK – which may not be the intended Station.
- If the password/PSKs are auto-generated, they may be cumbersome to enter by the user. If they're user-generated, it can become tedious – potentially leading to poorly chosen passwords.
- For a multi-station password/PSK credential, there's no reliable way to control which Station(s) are allowed to use the credential.
- Some devices use MAC address randomization prior to AP association for privacy protection. This can prevent the methods used for establishing initial MAC-to-PSK associations from working.

What is needed is a form of Station identification and credential that is nontransferable and attestable. In other words, Stations need an identifier that can be proven using secure methods and cannot be easily transferred from one Station to another. One such method will be discussed in Section 3.

## 2.2. WPA3-Personal

As discussed in Section 2.1.2, the basic messaging and authentication of WPA3-Enterprise is identical to WPA2-Enterprise – and the method to exchange the PMK is unchanged. WPA3-Personal, however, uses a dramatically different method than WPA2-Personal for mutually deriving the PMK. Rather than using the password/PSK (and SSID) to derive the PMK, WPA3-Personal uses an exchange called “Simultaneous Authentication of Equals” (SAE). SAE is a PAKE-based (Password Authenticated Key Exchange) cryptographic exchange that can derive a cryptographically strong shared secret from a low-entropy password.

This method solves one of the largest issues with WPA2-Personal shared passwords: the ability to derive and Station’s PTK and decrypt all the traffic on the network – including traffic captured before the PTK was known. Here’s an illustration of the SAE exchange:

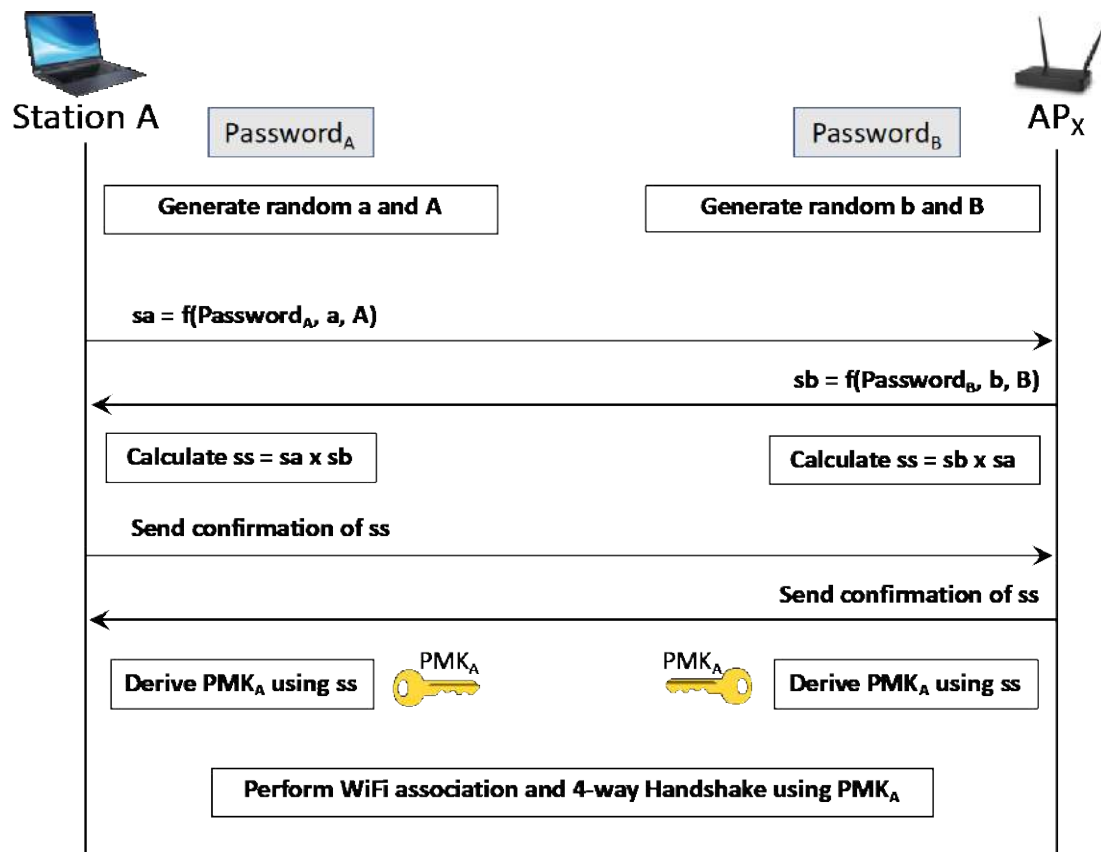


Figure 9: WPA3 Personal SAE Exchange

Some advantages of the SAE exchange:

- Both parties can calculate and send their components independently
- The computations used to derive the shared secret/PMK ensure that the password and the PMK cannot be determined
- Ensures perfect forward secrecy



### 3. Provisioning Station Credentials using DPP

The Wi-Fi Alliance Easy Connect™ specification defines mechanisms for provisioning a Station with Wi-Fi network credentials without user entry of passwords/keys or installation of enterprise certificates. While simplifying the task of device onboarding, Easy Connect also provides stronger security methods, per-device credentials, revocation, mutual authentication, and reliable/secure device- and group-level identification.

#### 3.1. DPP Provisioning

There are a variety of means in Easy Connect by which a device can be discovered and authenticated for provisioning. Figure 10 and Figure 11 illustrate the provisioning of a DPP “Connector” and PSK, respectively, on devices using a scanned QR code to provide network authentication and connectivity.

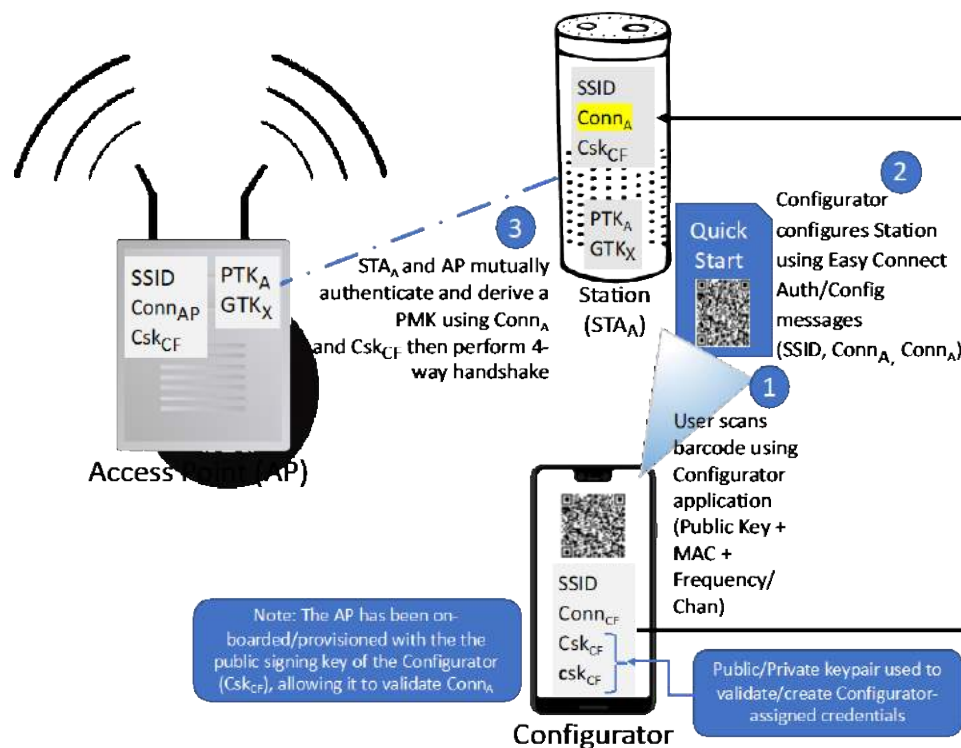
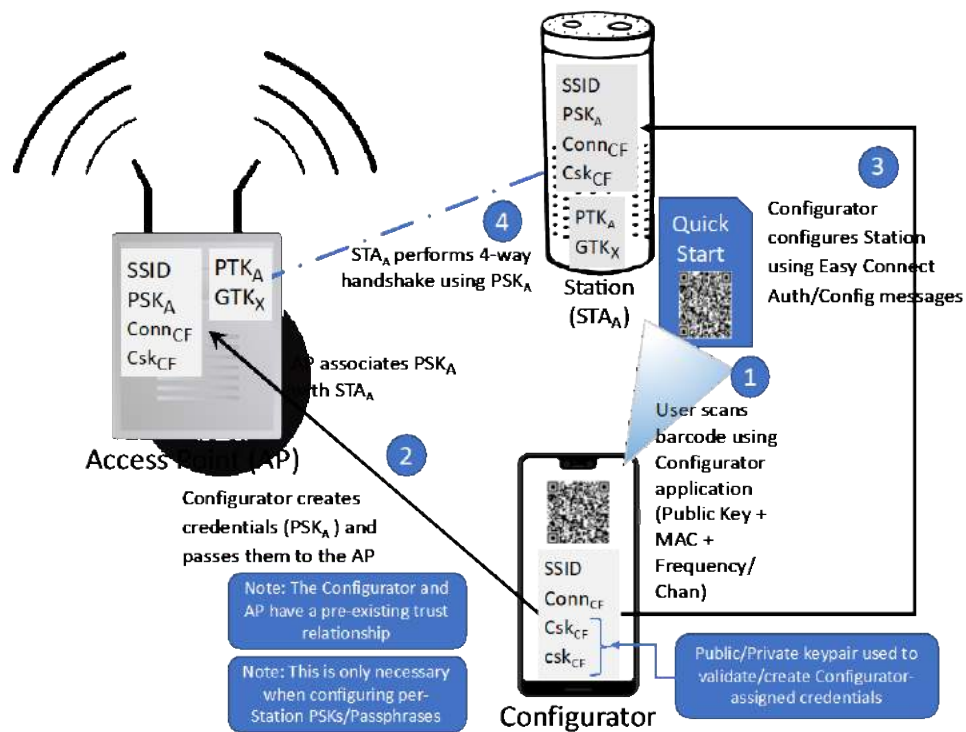


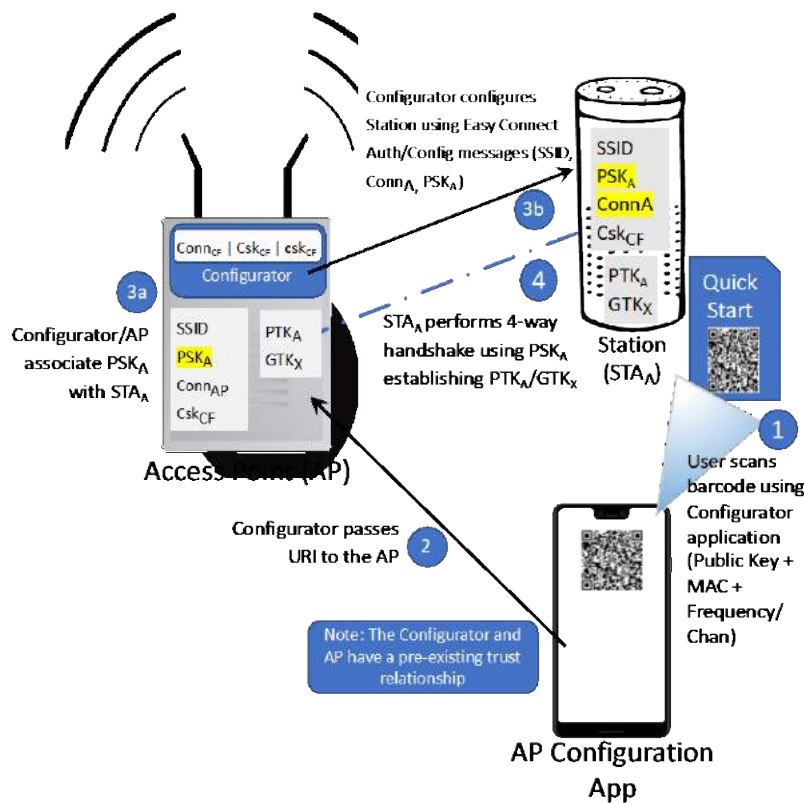
Figure 10: DPP Direct Provisioning of Connector Using QR Code





**Figure 11: DPP Direct Provisioning of PSK Using QR Code**

Figure 10 and Figure 11 illustrate a stand-alone Configurator directly provisioning the on-boarded Station (STA). But it can often be advantageous to have the Configurator co-located on the AP. This model ensures that (a) the Configurator is always accessible (which simplifies the discovery/initiation process) and (b) ensures that credentials stored by the Configurator are not easily lost (e.g. if the device with the Configurator is lost or the application is deleted). Figure 12 illustrates a Configurator which operates on the AP.

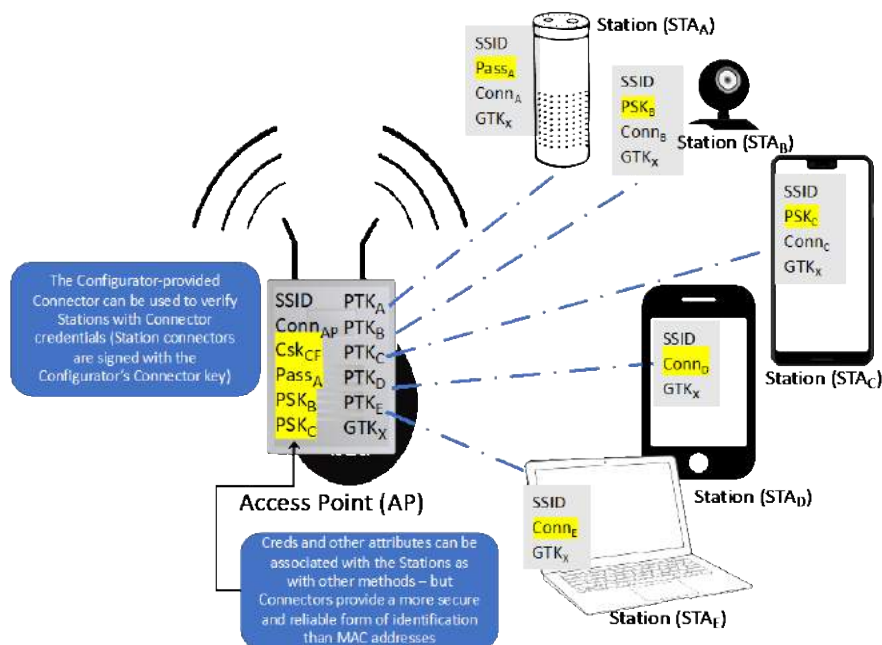


**Figure 12: DPP AP-based Provisioning of PSK using QR code**

Regardless of how provisioning is performed, all Stations can be provisioned with a PSK, WPA2 password, Enterprise credentials (certificate), WPA3/SAE Password, and/or DPP Connector credentials. The Configurator can offer more than one type of credential to a Station. And a Station can accept more than one credential. A network can be provisioned with a combination of credentials depending upon the

credential types supported on the AP, the credential types supported on the Stations, and what credentials the Configurator creates/supplies.

Below is an illustration showing Stations provisioned with a combination of different credential types:



**Figure 13: DPP Provisioning of PSKs**

The DPP onboarding process also provides the opportunity for the Configurator application (or its proxy) to interact with the user and configure the Station and the AP/Wi-Fi network according to user interaction/direction or other technologies. For example, during the onboarding process, the user may designate that a Station may only connect to the Internet and not to other devices on the home network, that a device should not be able to use more than a particular bitrate of uplink bandwidth, or that a Station should not be able to use the uplink during particular hours of the day

### 3.2. DPP/EasyConnect Connectors

Many of the solutions currently deployed to provide device/Station-level policy depend upon MAC addresses to identify devices and associate per-device and per-group policy, as described in section 2.1.3. There are a number of issues using MAC addresses as identifiers:

1. MAC addresses are easily observable via traffic monitoring – even by parties that are not part of the Wi-Fi network (since every Wi-Fi packet contains a MAC address),
2. The MAC address of an adapter is changeable. This is especially problematic when an AP/gateway with shared password credentials attempts to assign policy based on MAC address. It is trivial for one Station to take on another Station's MAC-associated policy by simply cloning the MAC address. E.g. A child in a household can circumvent MAC-associated network time or usage restrictions by stealing a parent's MAC address or assigning a random MAC address.
3. Due to (1), MAC addresses have become a privacy concern. Device manufacturers have announced initiatives to use randomized MAC addresses to avoid device/identity tracking.

Today's systems that utilize MAC-based policy assignment will likely encounter issues as these privacy initiatives are implemented. [4][5]

DPP solves the issue of identification with *Connectors*. Connectors provide a secure and durable identity for a Station while also enabling mutual authentication and verifiable metadata. Connectors are JSON Objects created by the Configurator at the time of onboarding, cryptographically signed using the private key of the Configurator ( $\text{esk}_{\text{CF}}$ ), and encoded using JWS (JSON Web Signature) serialization.

Here is an example of a Connector, prior to encoding and signing using JWS:

```
{
 "groups": [
 {
 "groupId": "home",
 "netRole": "sta"
 },
 {
 "groupId": "cottage",
 "netRole": "sta"
 }
],
 "expiry": "2019-01-31T22:00:00+02:00",
 "netAccessKey": {
 "kty": "EC",
 "y": "LUSDBmn7nv-LCnn6fBoXKsKpLGJiVpY_knTckGgsgeu",
 "x": "Xj-zV2iEiH8XwyA9ijpsL6xyLvDiIBthrHO8ZVxwmpA",
 "crv": "P-256"
 }
}
```

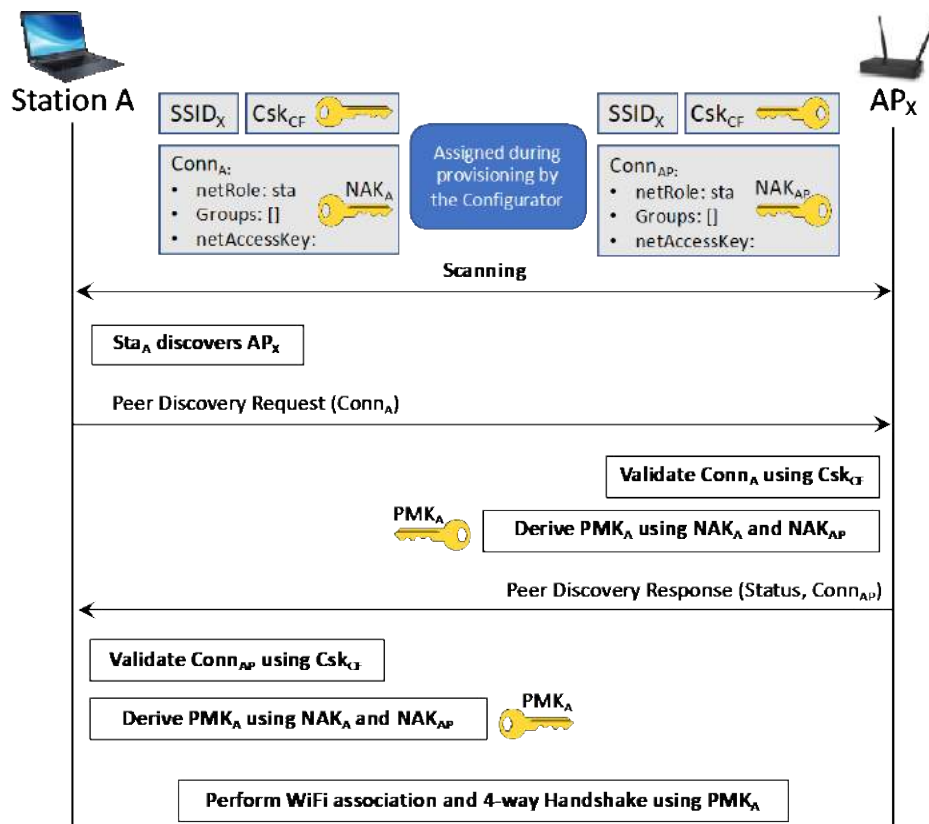
### Figure 14: DPP Connector Example

Once encoded and signed, the same Connector appears as three base64-encoded fields separated by “.” characters – with the first group being JWS “protected” header values, then the encoded Connector, followed by a signature of the two fields. The signature (or hash of the signature) can be used to identify the Connector. Here’s an example of a JWS-encoded and signed Connector:

eyJ0eXAiOiJkcHBDb24iLCJraWQwIj0rTWNlZ0RCUGlOWlZha0FzQlplPek9vq3N2UWprcl9uRUFWOXVGLUVebVZFtiwiYWxnIjoirVMMyNTYifQ  
.  
.eyJncm9lcHMiolt7Imdyb3VSWSQiOiJob211TiwbmV0Um9sZSI6InN0YSJ9LHsiZ3JvdXBjZCI6ImNvdHRhZ2UilCJuZXRsb2x  
lIjoic3RhInldLCJuZXRBRY2Nlc3NLZXkiOnsia3R5IjoirUMiLCJjcnyYiOiJQLTl1NiIsIngioiJYaill6vjUpRWllOFh3eUE5aW  
pwc0w2eHlMdkRpSUJ0aHJTtzhavnh3bXBBIiwieSI6IkxVc0RcbW43bnYtTENubjZmQm9YS3NLcExHSmlWcFlfa25UY2tHZ3NnZ  
VUifSwizXhwaXJ5IjoimjAxOS0zMWS0zMVQyMjowMDowMCswmjowMCJ9  
.  
8fJSNCpdjv5BEffmlqEbBNTAHz2L6c 22Uvr9KYjtAw88VfvEUWiruECUSJCUVFqvlyDEE4RJVDtIw3aUDhlMw

### Figure 15: Encoded DPP Connector Example

As part of the Wi-Fi authorization process, a Station supporting DPP can authenticate with the Access Point using the Connector supplied to the Station during provisioning. Figure 16 illustrates the mutual authentication and derivation of the PMK using Connectors.



**Figure 16: DPP Connector-based Authentication**

Note that Connector-based authentication addresses the major issues of WPA2-Personal described in Section 2.1.1 – by utilizing a PMK/PSK that’s device-specific and non-transferrable. And they also address the major issues of WPA2/3-Enterprise described in Section 2.1.2 – since Connectors are easily generated and easily validated without the need for enterprise infrastructure and provisioning.

### 3.3. Applications of DPP Connectors

As reliable identifiers, Connectors allow for the association of metadata with a Station. During the DPP device onboarding process, arbitrary metadata can be associated with a newly onboarded Station by associating the metadata with the Station’s connector (Specifically, the Connector signature or a hash of the signature).

For instance, during the provisioning process, a user could designate that a device should only operate during particular hours. Or a Station can provide, via the DPP Configuration Request object, a MUD (Manufacturer Usage Description) URL that – when processed – indicates that a Station/Device should only connect to a particular host. These provisioning attributes can be stored on the AP by associating them with the Station’s Connector.

Connectors also solve the issues related with MAC address reassignment described in Section 2.1.3. When the device authenticates – as shown in Figure 16 – the MAC address of the Station can be associated with the Station’s Connector after mutual authentication is completed. MAC-based policies/ACLs can be implemented on an AP/gateway without any requirement on MAC addresses being immutable across Station associations.

For example, the application of time-based Internet access restrictions for a Station “kidtab” for 10pm to 6am can be implemented by an AP/gateway using DPP Connectors via the following steps:

1. The Configurator onboarded “kidtab” and assigned it Connector C
2. The Configurator (or proxy application) is used to configure an access restriction for “kidtab” from 10pm to 6am
3. The Configurator adds a time-based policy attribute for 10pm to 6am (daily) and associates it with Connector C on the AP/gateway (using previously established credentials).
4. When “kidtab” authenticates/associates with the network using Connector C, the AP/gateway associates the MAC address with Connector C (the “kidtab” Connector).
5. Once “kidtab” is fully connected, the AP/gateway looks up the policies for “kidtab”, finds an associated time-based policy, and sets a timer for 10pm.
6. When the timer goes off at 10pm, the AP/gateway sets packet blocking rules that prevent any traffic with the source MAC address associated with “kidtab” (determined in step 3) and a destination IP outside the local network.

With MAC-associated policy (as is the case with Wi-Fi policy systems today) steps (3) and (4) can fail if/when endpoints utilize MAC randomization – even when per-Station passwords are used. And when shared passwords are used, these kinds of policies can be easily defeated by simply cloning the MAC address of another authorized Station.

Another application of DPP Connectors is reprovisioning. Since all DPP-provisioned Stations and APs are configured with both a Connector and the public key of its Configurator (enabling mutual authentication), the Stations and/or AP can be dynamically and securely reprovisioned at any time by the Configurator using the Connector-based trust relationship. Reconfiguration can be used to update the network configuration of a Station without user intervention – for example to change the SSID of the network. DPP reconfiguration can also be used to provision new Connectors (with new attributes) based on changes made in the Configurator application.

## 4. Wireless Network Segmentation

The number and nature of wireless devices connected to home/SOHO networks has changed dramatically since Wi-Fi was introduced. Smart TVs/speakers, set-top boxes, and streaming devices interact with both the local area network and cloud-based services, while IoT devices such as doorbells, lightbulbs, thermostats, security cameras, and appliances interact and integrate exclusively with cloud-based services.

Additionally, the single- or limited-purpose nature of IoT devices makes them more commoditizable. For instance, to many people, most brands of smart plugs, switches, and bulbs are inconsequential. The considerations are (a) are they compatible with the smart speaker/home automation system that people use and recognize (e.g. “Alexa-Compatible”, “Google Home Compatible”, “Apple HomeKit Compatible”, etc.) and (b) purchase cost. Each device – and its associated (and necessary) cloud-based controller – represent a distinct risk to the user’s network and even to networks/systems outside the home. One compromised cloud-based controller can be used to facilitate a botnet-based DDoS (Distributed Denial of Service) attack.

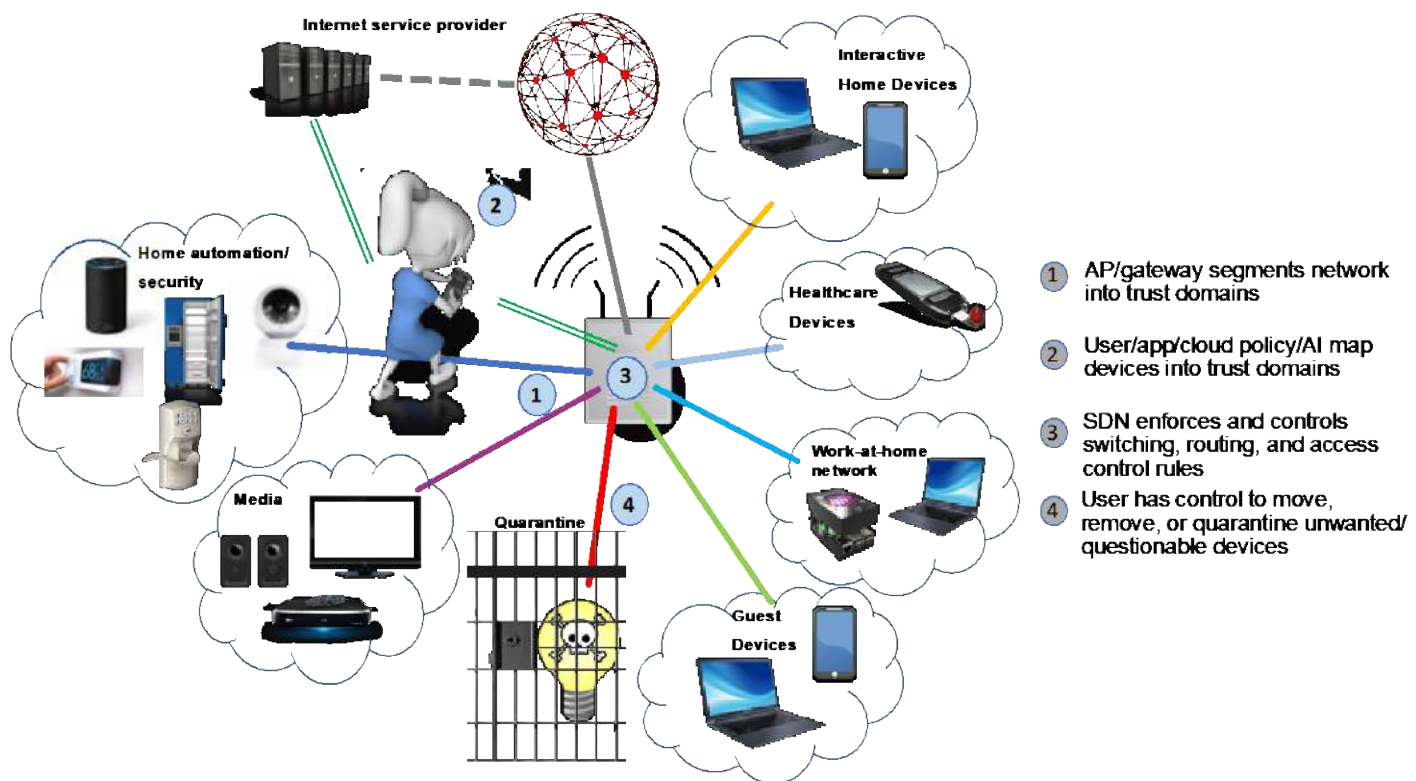
With more devices there’s also more demand for shared resources. Most notably, Wi-Fi bandwidth and Internet uplink bandwidth can become strained. And there may be a desire to prioritize traffic to reduce latency for applications such as interactive gaming and video conferencing.

One solution to address the current needs of the modern Wi-Fi network and devices – and to provide capabilities for future network requirements – is *network segmentation*. Network segmentation can be implemented in AP/gateways using a combination of virtual bridges and SDN (software-defined networking). But this form of network segmentation can only be reliably and securely implemented when there's a solid means of attestation for a device's identity and layer 2 (MAC) addresses – as described in Section 3.

Network segmentation can enable:

- The separation of Stations by risk category and/or function. E.g. Home automation (HA)/security devices can be allowed to connect to the Internet but cannot initiate connections to devices in other segments. Or devices with a higher quality of credential (e.g. a DPP Connector) can be separated from those with lower-quality credentials (e.g. a PSK).
- Segment-level access controls. Stations in some segments can discover/initiate connections to devices in other segments, but not vice-versa. E.g. Interactive devices such as smart phones and tablets can initiate connections with home automation/security devices, but those HA/security devices cannot access non-HA/security devices
- The “quarantine” of Stations that are suspected of being compromised – either explicitly via user direction or automatically, based on behavior/AI or external threat notifications
- The application segment-level policies – such as usage and time quotas.

Figure 17 illustrates an example network segmentation.



**Figure 17: Example Network Segmentation**

## 4.1. Network Segmentation using DPP Connectors

The disposition of Stations into particular segments can be done initially during onboarding – based on Station information provided in the DPP Config Request, initial or historical profiling of the Station (e.g. via device fingerprinting), and/or user direction (e.g. the user choosing a segment from a drop-down on the Configurator UI). A Station can also be moved from one segment to another, based on preference changes or updated information. E.g. If an advisory is published about a particular device, the AP/Gateway could move the Station into a “Quarantine” segment until its software is updated.

To implement segmentation using SDN an AP/Gateway minimally requires:

1. Each Station to have a unique identity – which allows a Station to be associated with a segment.
2. Each Station to have its own security credentials which are associated with the identity and non-transferrable.
3. The MAC address of each Station – which allows SDN rules to be applied to the Station’s traffic.

As outlined in Section 3.2, DPP provides (1) and (2) directly. And as described in Section 3.3, it indirectly provides (3) in a way that is much more robust than today’s methods that presume immutable MAC addresses.

For example, the provisioning of Station “kidtab” into the network segment “Home-Secure” (a segment designated for laptops, tablets, and smartphones for household residents) could be accomplished with the following steps:

1. The Configurator (or proxy application) is used to onboard “kidtab”. The Configurator creates Connector C for it with a “groups” containing “net-seg” and provides Connector C and an SSID of “HomeNet” in the DPP exchange.
2. The Configurator adds a “net-seg” attribute with the value “Home-Secure” and associates the attribute with the Connector C on the AP/Gateway (using previously established credentials)
3. When “kidtab” mutually authenticates with the “HomeNet” AP/gateway using Connector C, the AP/gateway associates the MAC address in the message exchange with “kidtab”.
4. Once “kidtab” is fully associated, the AP/gateway looks up the “net-seg” policy for “kidtab”, finds the name “Home-Secure” and:
  - a. allocates an IP address for “kidtab” in an IP subnet associated with the “Home-Secure” network segment.
  - b. Ensures that SDN rules are written appropriately for “kidtab” to be in the “Home-Secure” segment – and to enforce any segment- or Station-specific policy – using the MAC address for “kidtab” (determined in Step 3).
5. If/when “kidtab” disassociates from “HomeNet”, the AP/gateway removes the MAC association for “kidtab”, the IP address assignment, and all SDN rules associated with its MAC address. This effectively removes it from the “Home-Secure” network.

There are, of course, many ways to implement segmentation. But invariably there will be both Layer 2 and Layer 3 processing required to implement segmentation – especially for Station-specific rule enforcement.

## 5. Conclusion

The amazing success of Wi-Fi now presents its greatest challenge: Supporting the security, scaling, and privacy needs that are required going forward while still supporting the successful Wi-Fi technologies that we’re all using today.



DPP bridges that gap by supporting the provisioning of current WPA2 and WPA3 credentials while also supporting advanced credentials that allow for mutual authentication and identification. And by integrating with the existing Wi-Fi 4-way handshake protocol, DPP is able to provide that support without introducing new hardware requirements on APs or Stations.

The features of DPP in and of themselves make a compelling case for the protocol. But when you take into consideration the security issues with IoT devices, the need to provide per-Station policies for managing/prioritizing access for non-IoT devices, and the privacy concerns related to MAC addresses, the importance and value of having a non-spoofable Station and AP identity becomes obvious. And the application of per-Station policies and network segmentation – built on the foundation of DPP – can ensure that Wi-Fi will be able to meet the evolving needs of users and devices.

## Abbreviations

|        |                                                   |
|--------|---------------------------------------------------|
| AAA    | Authentication Authorization and Accounting       |
| ACL    | Access control list                               |
| AP     | Access point                                      |
| bps    | Bits per second                                   |
| DDoS   | Distributed denial of service                     |
| DPP    | Device provisioning protocol                      |
| EAP    | Extensible authentication protocol                |
| EAPoL  | Extensible authentication protocol (EAP) over LAN |
| GTK    | Groupwise temporal key                            |
| IEEE   | Institute of Electrical and Electronics Engineers |
| JWS    | JSON Web Signature                                |
| MAC    | media access control                              |
| MUD    | Manufacturer usage description                    |
| PAKE   | Password authenticated key exchange               |
| PMK    | Pairwise master key                               |
| PSK    | Pre-shared key                                    |
| PTK    | Pairwise temporal key                             |
| QR     | Quick Response (code)                             |
| RADIUS | Remote Authentication Dial-In User Service        |
| SAE    | Simultaneous authentication of equals             |
| SCTE   | Society of Cable Telecommunications Engineers     |
| SDN    | Software Defined Networking                       |
| SOHO   | Small office home office                          |
| SSID   | Service Set Identifier                            |
| VLAN   | Virtual local area network                        |
| WEP    | Wired Equivalent Privacy                          |
| WPA    | Wi-Fi Protected Access                            |
| WPA2   | Wi-Fi Protected Access 2                          |

## Bibliography & References

- [1] Device Provisioning Protocol Version 1.2; Wi-Fi Alliance
- [2] WPA3™ Specification Version 2.0; Wi-Fi Alliance

- [3] IEEE 802.11i-2004: *IEEE Standard for information technology-Telecommunications and information exchange between systems-Local and metropolitan area network - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications*; Institute of Electrical and Electronics Engineers
- [4] Apple Knowledge Base Article 211227: *Use private Wi-Fi addresses in iOS 14, iPadOS 14, and watchOS 7*; <https://support.apple.com/en-us/HT211227>
- [5] Android 10 Privacy and location/Privacy Changes: *MAC Address Randomization*; <https://developer.android.com/about/versions/10/privacy/changes#randomized-mac-addresses>

# Enforcing Social Distancing Using Computer Vision and Deep Learning

A Technical Paper prepared for SCTE•ISBE by

**Wael Guibene**

Director – Wireless R&D  
Charter Communications  
6360 S Fiddlers Green Cir, Englewood, CO 80111  
Wael.Guibene@charter.com

**Hossam Hmimy**

Sr. Director – Wireless R&D  
Charter Communications  
6360 S Fiddlers Green Cir, Englewood, CO 80111  
Hossam.Hmimy@charter.com

# Table of Contents

| <b>Title</b>                                                        | <b>Page Number</b> |
|---------------------------------------------------------------------|--------------------|
| 1. Introduction.....                                                | 3                  |
| 2. Computer Vision and DL-based Social Distancing Application ..... | 3                  |
| 3. System Description .....                                         | 4                  |
| 4. People Detection Algorithm.....                                  | 5                  |
| 5. People Tracking via Centroid .....                               | 6                  |
| 6. Social Distancing PoC.....                                       | 8                  |
| 7. Conclusion.....                                                  | 10                 |
| Bibliography & References.....                                      | 11                 |

## List of Figures

| <b>Title</b>                                                                                                       | <b>Page Number</b> |
|--------------------------------------------------------------------------------------------------------------------|--------------------|
| Figure 1 Performance overview of the most popular object detection models on PASCAL-VOC and MS-COCO datasets ..... | 4                  |
| Figure 2. E2E Detection and Tracking Flow. ....                                                                    | 4                  |
| Figure 3 Classification using YOLO v3 Framework. ....                                                              | 5                  |
| Figure 4. Centroid tracking step 1. ....                                                                           | 6                  |
| Figure 5. Centroid tracking step 2. ....                                                                           | 7                  |
| Figure 6. Centroid tracking step 3. ....                                                                           | 7                  |
| Figure 7 Centroid tracking step 4 .....                                                                            | 8                  |
| Figure 8. Social Distance fully respected. ....                                                                    | 9                  |
| Figure 9. Social Distance not respected. ....                                                                      | 9                  |
| Figure 10. Social Distance partially respected.....                                                                | 10                 |

# 1. Introduction

A defining moment of the century, so far, is the unprecedented impact that COVID-19 has brought to the economies world-wide, the populations and defining new norms in society.

Our paper details how we can enforce the new rules of society like social distancing and wearing face masks in open-spaces using computer vision and deep learning through:

- Detecting people on a particular scene,
- Calculating and monitoring the distances between the different people,
- Tracking movements and segregating moving people (might come close to each other during brief moments) from people standing still and violating the social distancing rules, and
- Creating alerts (audio, visual, light...) to enforce social distancing.

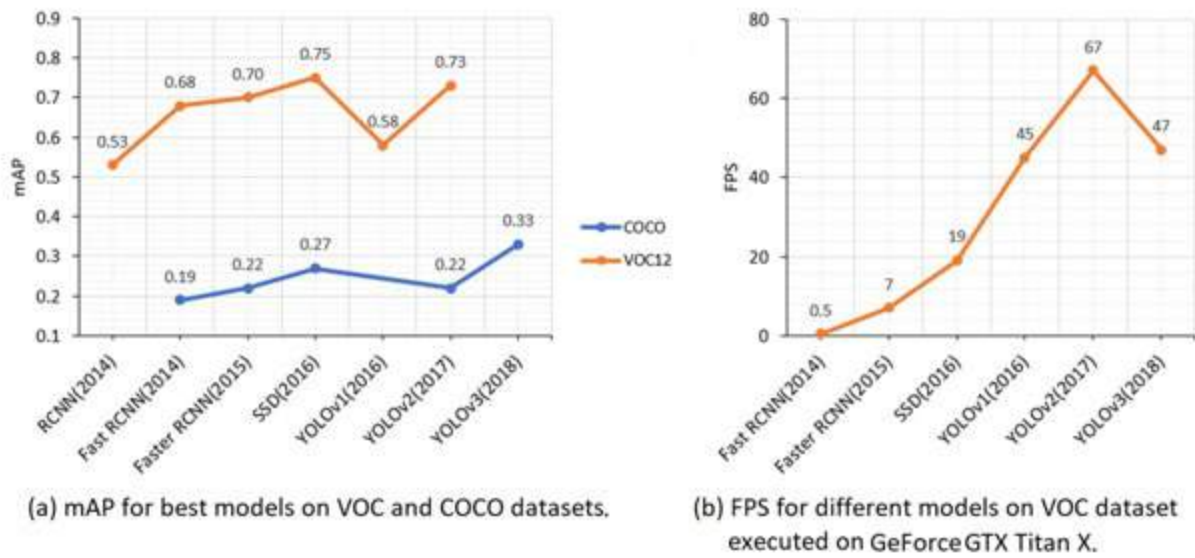
Our approach also ensures that, in confined spaces, we count people and create visual and audio alerts when the number of people exceeds the Center for Disease Control (CDC) guidelines (10 people per room).

We detail in the paper the computer vision and deep learning frameworks we used to achieve high confidence in detecting human presence, calculating and calibrating distances in the frames, and removing false positives (eg. people crossing paths while walking versus people standing still).

## 2. Computer Vision and DL-based Social Distancing Application

The emergence of deep learning has brought the best performing techniques for a wide variety of tasks and challenges including medical diagnosis, machine translation, speech recognition, and a lot more. Most of these tasks are centered around object classification, detection, segmentation, tracking, and recognition.

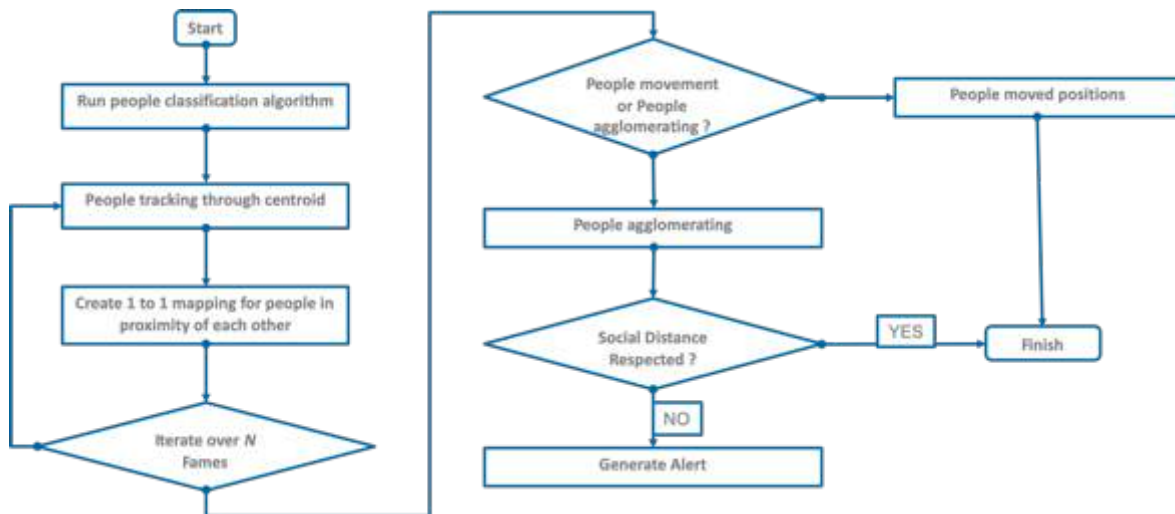
In recent years, the convolution neural network (CNN) based architectures have shown significant performance improvements that are leading towards the high quality of object detection, as shown in Fig. 1, which presents the performance of such models in terms of mAP and FPS on standard benchmark datasets, PASCAL-VOC and MS-COCO, and similar hardware resources. In this paper, a deep learning-based framework is proposed that utilizes object detection and tracking models to aid in the social distancing remedy for dealing with the escalation of COVID-19 cases. To maintain the balance of speed and accuracy, YOLO v3 alongside Centroid tracking are utilized as object detection and tracking approaches while surrounding each detected object with bounding boxes. Later, these bounding boxes are utilized to compute the pairwise L2 norm with computationally efficient vectorized representation for identifying the clusters of people not obeying the order of social distancing. Furthermore, to visualize the clusters in the live stream, each bounding box is color-coded based on its association with the group where people belonging to the same group are represented with the same color. Each surveillance frame is also accompanied with the streamline plot depicting the statistical count of the number of social groups and an index term (violation index) representing the ratio of the number of people to the number of groups. Furthermore, estimated violations can be computed by multiplying the violation index with the total number of social groups.



**Figure 1 Performance overview of the most popular object detection models on PASCAL-VOC and MS-COCO datasets**

### 3. System Description

In this section, we describe the end-to-end (E2E) system flow for people detection, tracking and social distance measurement.



**Figure 2. E2E Detection and Tracking Flow.**

The flow as depicted in Fig.2 starts when the system is receiving frames via RTSP from a networked camera (WiFi, Ethernet, CBRS, LTE...). The algorithm performs the object classification until classifying objects as people. Each person in the scene is tracked via centroid algorithm. A “security/privacy zone” surrounding each person of X ft is created and the algorithms keeps tracking distances between the closet centroids to each other. In order to minimize false alarm, we re-iterate over N frames in order to rule out people moving from people standing still. If the same centroids are identified in a single group over the N frames, the algorithm classifies the group as standing still and conglomerating, if the group is not respecting the social distance of X an alert (visual, light, voice) is issued to remind the group of social distancing rules.

## 4. People Detection Algorithm

From the Camera stream, we run an object identification and classification algorithm to infer with high precision the presence of people in the video stream.

Each frame is sub-divided into smaller Regions of Interests (ROIs). Each boundary box or ROI contains 5 elements: (x, y, w, h) and a box confidence score. The confidence score reflects how likely the box contains an object (objectness) and how accurate is the boundary box. We normalize the bounding box width w and height h by the image width and height. x and y are offsets to the corresponding cell. Hence, x, y, w and h are all between 0 and 1. Each cell has 20 conditional class probabilities. The conditional class probability is the probability that the detected object belongs to a particular class (one probability per category for each cell). So, our approach’s prediction has a shape of  $(S, S, B \times 5 + C) = (7, 7, 2 \times 5 + 20) = (7, 7, 30)$ . The major concept is to build a CNN to predict a  $(7, 7, 30)$  tensor. It uses a CNN to reduce the spatial dimension to  $7 \times 7$  with 1024 output channels at each location. The algorithm performs a linear regression using two fully connected layers to make  $7 \times 7 \times 2$  boundary box predictions. To make a final prediction, we keep those with high box confidence scores (greater than 25%) as our final predictions.




Figure 3 Classification using YOLO v3 Framework.

## 5. People Tracking via Centroid

### Step 1: Accept bounding box coordinates and compute centroids

The centroid tracking algorithm assumes that we are passing in a set of bounding box (x,y)-coordinates for each detected object in every single frame. These bounding boxes can be produced by our object detector provided that they are computed for every frame in the video. Once we have the bounding box coordinates we must compute the “centroid”, or more simply, the center (x,y)-coordinates of the bounding box.

Figure 4 below demonstrates accepting a set of bounding box coordinates and computing the centroid. Since these are the first initial set of bounding boxes presented to our algorithm we will assign them unique IDs.

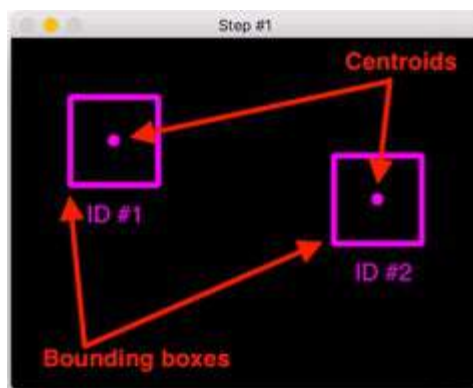



Figure 4. Centroid tracking step 1.

### Step 2: Compute Euclidean distance between new bounding boxes and existing objects.

For every subsequent frame in our video stream we apply Step #1 of computing object centroids; however, instead of assigning a new unique ID to each detected object (which would defeat the purpose of object tracking), we first need to determine if we can associate the new object centroids (yellow) with the old object centroids (purple). To accomplish this process, we compute the Euclidean distance (highlighted with green arrows) between each pair of existing object centroids and input object centroids. From figure 5, we can see that we have this time detected three objects in our image. The two pairs that are close together are two existing objects. We then compute the Euclidean distances between each pair of original centroids (yellow) and new centroids (purple). But how do we use the Euclidean distances between these points to actually match them and associate them? → Step 3






**Figure 5. Centroid tracking step 2.**

### **Step 3: Update (x, y)-coordinates of existing objects**

The primary assumption of the centroid tracking algorithm is that a given object will potentially move between subsequent frames, but the distance between the centroids for frames  $F_t$  and  $F_{(t+1)}$  will be smaller than all other distances between objects. Therefore, if we choose to associate centroids with minimum distances between subsequent frames we can build our object tracker. In figure below, we can see how our centroid tracker algorithm chooses to associate centroids that minimize their respective Euclidean distances.

The new point that appears at the bottom left mean that we see a new object we need to register → Step 4




**Figure 6. Centroid tracking step 3.**

### **Step 4: Register new objects**

In the event that there are more input detections than existing objects being tracked, we need to register the new object. “Registering” simply means that we are adding the new object to our list of tracked objects by:

- Assigning it a new object ID, and
- Storing the centroid of the bounding box coordinates for that object, then
- We can then go back to Step #2 and repeat the pipeline of steps for every frame in our video stream.

Figure 7 depicts the process of using the minimum Euclidean distances to associate existing object IDs and then registering a new object.



**Figure 7 Centroid tracking step 4**


#### **Step 5: Deregister old objects**

We deregister old objects when they cannot be matched to any existing objects for a total of N subsequent frames.


## **6. Social Distancing PoC**

In order to validate our approach, we implemented and validated our algorithms in our offices in Greenwood Village.


Bellow figures depict the outcome of the PoC:



**Figure 8. Social Distance fully respected.**



**Figure 9. Social Distance not respected.**



**Figure 10. Social Distance partially respected.**

## 7. Conclusion

In this paper, we presented a novel approach to an AI-assisted social distancing application for safer back to work situations.

This solution can be applied to indoor or outdoor scenarios ensuring social distances are met and respected.

We have deployed and tested our solution in our lab, and it has shown promising results and good accuracy for measuring distances and alerting (via sound and lights) when social distancing is not respected.

## Bibliography & References

- [1] S. A. Niyogi and E. H. Adelson, “Analyzing gait with spatiotemporal surfaces,” in Proceedings of 1994 IEEE Workshop on Motion of Nonrigid and Articulated Objects. IEEE, 1994, pp. 64–69.
- [2] Z.-Q. Zhao, P. Zheng, S.-t. Xu, and X. Wu, “Object detection with deep learning: A review,” IEEE transactions on neural networks and learning systems, vol. 30, no. 11, pp. 3212–3232, 2019.
- [3] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in Advances in neural information processing systems, 2012, pp. 1097–1105.
- [4] S. Ren, K. He, R. Girshick, and J. Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” in Advances in neural information processing systems, 2015, pp. 91–99.
- [5] X. Chen and A. Gupta, “An implementation of faster rcnn with study for region sampling,” arXiv preprint arXiv:1702.02138, 2017.
- [6] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” in Proceedings of the IEEE 10 conference on computer vision and pattern recognition, 2016, pp. 779–788.
- [7] M. Putra, Z. Yussof, K. Lim, and S. Salim, “Convolutional neural network for person and car detection using yolo framework,” Journal of Telecommunication, Electronic and Computer Engineering (JTEC), vol. 10, no. 1-7, pp. 67–71, 2018.
- [8] R. Eshel and Y. Moses, “Homography based multiple camera detection and tracking of people in a dense crowd,” in 2008 IEEE Conference on Computer Vision and Pattern Recognition. IEEE, 2008, pp. 1–8.
- [9] D.-Y. Chen, C.-W. Su, Y.-C. Zeng, S.-W. Sun, W.-R. Lai, and H.- Y. M. Liao, “An online people counting system for electronic advertising machines,” in 2009 IEEE International Conference on Multimedia and Expo. IEEE, 2009, pp. 1262–1265.
- [10] C.-W. Su, H.-Y. M. Liao, and H.-R. Tyan, “A vision-based people counting approach based on the symmetry measure,” in 2009 IEEE International Symposium on Circuits and Systems. IEEE, 2009, pp. 2617–2620.
- [11] J. Yao and J.-M. Odobez, “Fast human detection from joint appearance and foreground feature subset covariances,” Computer Vision and Image Understanding, vol. 115, no. 10, pp. 1414–1426, 2011.
- [12] B. Wu and R. Nevatia, “Detection and tracking of multiple, partially occluded humans by bayesian combination of edgelet based part detectors,” International Journal of Computer Vision, vol. 75, no. 2, pp. 247–266, 2007.

# **Access Capacity Planning: Staying Well Ahead Of Customer Demand Helped Ensure Stability During COVID-19**

A Technical Paper prepared for SCTE•ISBE by

**Bruce E. Barker Jr.**

Vice President, Capacity Engineering  
Next Generation Access Network (NGAN), Comcast Cable  
1701 JFK Blvd., Philadelphia, PA 19103  
609-685-4269  
Bruce\_Barker@cable.comcast.com

**Claude Bou Abboud**, Sr. Director, Comcast Cable

**Erik Neeld**, Sr. Director, Comcast Cable

# Table of Contents

| Title                                | Page Number |
|--------------------------------------|-------------|
| 1. Introduction.....                 | 3           |
| 2. Historic Context.....             | 3           |
| 3. Problem Statement .....           | 4           |
| 4. Solutions.....                    | 6           |
| 4.1 Proactive Capacity Planning..... | 7           |
| 4.2 New and Modified Reporting ..... | 9           |
| 4.3 Capacity Acceleration .....      | 10          |
| 5. Conclusion .....                  | 14          |
| Acknowledgements .....               | 16          |
| Abbreviations.....                   | 17          |

## List of Figures

| Title                                                                                                                                                       | Page Number |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| Figure 1 – Traffic Growth Since Late February 2020 ( <i>Source: Comcast Capacity Management</i> ).....                                                      | 5           |
| Figure 2 – CMTS Upstream Service Group Utilization Shift ( <i>Source: Comcast Capacity Management</i> ) .....                                               | 5           |
| Figure 3 – Upstream Hourly Consumption Shift ( <i>Source: Comcast Capacity Management</i> ) .....                                                           | 6           |
| Figure 4 – Traffic Forecasting Using ETS Methodology ( <i>Source: Comcast Capacity Management</i> ) .....                                                   | 7           |
| Figure 5 – Comparison of Capacity Solutions ( <i>Source: Comcast Capacity Management</i> ) .....                                                            | 8           |
| Figure 6 – Example CMTS Scorecard ( <i>Source: Comcast Capacity Management</i> ) .....                                                                      | 9           |
| Figure 7 – Example CMTS SG Saturation or "Flatlining" ( <i>Source: Comcast Service Performance Database</i> ) .....                                         | 10          |
| Figure 8 – Example DOCSIS Upstream Channel Lineup With Fifth and Sixth Upstream Channel ( <i>Source: Comcast Next Generation Access Engineering</i> ) ..... | 11          |
| Figure 9 – Octave High Level System Diagram ( <i>Source: Comcast Next Generation Access Engineering</i> ) .....                                             | 12          |
| Figure 10 – Highly Utilized CMTS SG Trending ( <i>Source: Comcast Capacity Management</i> ) .....                                                           | 14          |

## 1. Introduction

As means to control the spread of the COVID-19 virus, many countries implemented measures to quarantine large portions of their populations in early 2020. Remote working and schooling policies were applied to a fairly large extent. Large portions of the workforce, especially in the hospitality industry, were laid off. Essentially, many people found themselves stuck at home for one reason or another.

This resulted in a significant increase in network traffic in many areas for many service providers. This was especially true for upstream traffic generated by various teleconferencing and VoIP-type protocols that were heavily leveraged to support remote working and schooling. Additionally, there was a noted shift in the daily peak time - from evenings to daytime in many locations.

Traffic increases caused an increase in the utilization of all network elements including cable modem termination system (CMTS) service groups (SGs) which provide connectivity between the core network the access network where customer homes and businesses reside. High utilization can result in potential impact to customer experience such as slow speeds, poor video quality, and intermittent service if not resolved. Speed test data and various industry models have helped us to define highly utilized CMTS SGs as those with utilization exceeding 85% consistently.

Sharp increases in traffic and CMTS SG utilization like what was observed during the COVID-19 event can be well-managed through a combination of proactive and reactive capacity management activities. Proactive measures are mainly focused on predictive traffic forecasting and capacity deployment planning. Reactive measures include a plethora of specialized reporting mechanisms, to help prioritize near-term capacity work, and adaptations to existing long-term plans where applicable.

Accelerated development of new solutions, techniques, and technologies is also a key aspect of any strategy during a crisis. To expand upstream capacity, development and testing for additional upstream channels was prioritized and successfully deployed. The profile management application (PMA) platform, known by the name “Octave” at Comcast and initially developed for DOCSIS 3.1 downstream capacity, was rapidly modified to enable more DOCSIS 3.0 upstream capacity. To help operationalize and deploy new solutions, other tools and automation were also quickly developed and modified.

Furthermore, the partnership between capacity management, engineering, finance, and technical operations teams on the planning and execution of capacity solutions was critical to tackling the challenge imposed by the sudden network growth. It was only through this partnership and alignment that extraordinary innovations and accomplishments were made. Node split work to mitigate highly utilized CMTS SGs, for example, increased 100% from normal levels as all teams rallied around a common goal. This would not have been possible otherwise. All efforts to manage capacity in extreme circumstances like during COVID-19 are only effective when there is good partnership and solid alignment across the organization.

## 2. Historic Context

There are various factors involved in the management of traffic growth and customer experience. Subscriber growth and increased usage per subscriber are key factors in driving traffic growth over time. In recent years, annual upstream network traffic growth has been in the 15-25% range, on average. Downstream traffic growth has typically been in the 25-35% range on average. Not accounting for seasonality, this has translated to averaged monthly growth of approximately 1-2% for upstream traffic and 2-3% for downstream traffic.



While traffic grows on some average amount, it typically does not grow in a linear or consistent manner in most locations. In reality, traffic follows seasonal patterns in which it grows faster or slower based on the time of year. Generally, traffic growth is higher during the autumn and winter months in most areas in the United States especially in the northern latitudes for a multitude of reasons. Schooling, for example, typically resumes in the late August and September timeframes with ever increasing reliance on the internet to complete homework assignments and projects. In this same timeframe, new primetime television shows and series are released both on traditional broadcast and over-the-top (OTT) providers. Not surprisingly, content from traditional broadcast providers is mostly watched via video-on-demand (VOD) platforms to better fit with hectic schedules. All OTT content and much of the VOD content is IP-based and delivered over DOCSIS for cable operators. Near the Thanksgiving and Christmas holidays, even more video content is made available as well as new games and gaming updates, all of which are delivered over DOCSIS and manifest in higher traffic than normal. And of course, winter weather events and generally unfavorable outdoor weather conditions cause many people to stay inside and spend their time watching more video content or playing more online games. Traffic then typically begins to flatten and even decline going into the spring and summer months, as people spend more time outdoors and new video content is limited.

Vacation areas also demonstrate seasonality on network traffic. Summer vacation homes near beaches, lakes, or mountains drive a significant amount of localized traffic, mostly during the summer months. Likewise, winter vacation homes near winter activity areas and in warmer southern latitudes (i.e. “snow-bird” communities) drive higher localized traffic during the winter months.

Daily peak usage time is another factor that may affect how capacity is managed. In most areas, the daily peak usage time happens during the evening hours. Some areas with a heavy penetration of business services customers may peak during the day.


Traffic growth and timing of traffic growth impacts the utilization of network elements including CMTS SGs and potentially customer experience. Most providers strive to account for the various factors related to traffic growth in order to minimize high utilization and ensure good product quality and service reliability at all times. At Comcast, the number of highly utilized CMTS SGs has typically been managed to a nominal number, even accounting for factors like seasonality and peak time usage.

There are various approaches to measuring CMTS SG utilization used across the industry. Our approach is based on calculating the 98<sup>th</sup> percentile of peak, of weekly 5-minute polling data averaged over four weeks. Each week, there are 2,016 5-minute utilization measurements collected. The data set is then sorted in descending order and the top 2% or 40 of these data points are removed. These top 2% or 40 data points are considered outliers. The 41<sup>st</sup> data point becomes the reported weekly utilization value for a given CMTS SG. We then calculate the average of the trailing four weeks to derive the official utilization value per CMTS SG. CMTS SGs with 4-week utilization values exceeding 85% are considered highly utilized. Normally, less than .1% of downstream SGs and less than .5% of upstream SGs exceed 85% utilization.

### **3. Problem Statement**


Rapid growth in network traffic can cause an increase in highly utilized CMTS SGs, resulting in potential impact to customer experience in some areas. In April 2020, upstream traffic increases in excess of 30% and downstream traffic increases in excess of 20%, on average, were not uncommon (Figure 1). This was especially true in more urbanized areas and areas where large populations of the workforce teleworked. Voice over IP (VoIP) and video teleconferencing traffic generated by remote working and

schooling applications increased in excess of 200% in many locations. Not surprisingly with limited entertainment options outside the home, gaming downloads increased 20-80% and streaming and web video increased 40% on average.




**Figure 1 – Traffic Growth Since Late February 2020** (Source: Comcast Capacity Management)

This rapid increase in upstream traffic manifested itself as an overall increase in upstream CMTS SG utilization, including an increase in highly utilized CMTS SGs (Figure 2).



**Figure 2 – CMTS Upstream Service Group Utilization Shift** (Source: Comcast Capacity Management)

Additionally, upstream peak traffic tended to shift from 9pm to 2pm local time. The highest hourly increase, of over 100%, was at 1pm (Figure 3). The traffic also flattened throughout the day and evening time from 10am to 10pm local. Most of the increase in daytime usage was, not surprisingly, driven by specific applications and protocols supporting virtual private network (VPN), VoIP, and video conferencing.



**Figure 3 – Upstream Hourly Consumption Shift** (Source: Comcast Capacity Management)

## 4. Solutions


The good news is, as part of an overall capacity engineering and management program, a combination of proactive and reactive measures can successfully mitigate and resolve the impacts of a rapid increase in network traffic due to an event like COVID-19, in a fiscally responsible manner. Proactive measures include predictive traffic forecasting, capacity planning, and proactive capacity deployment. Reactive measures include new and modified reporting, re-prioritization, and accelerated solution development/deployments.

Arguably, proactive capacity management efforts, and particularly those related to predictive forecasting and planning, set the foundation for all other efforts related to capacity management - whether during normal times or during events with extreme growth. According to Elad Nafshi, SVP of Next Generation Access Networks in Comcast, “proactive capacity management was the key to Comcast’s access network not only surviving but thriving during the COVID-19 event in the spring of 2020.” Having proactive plans already in place (and, in some cases, already in motion) made it much easier to enable reactive efforts like accelerating work and re-prioritizing focus. By the early summer, most highly utilized CMTS SGs from the spring were either resolved or had firm completion dates planned, near-term.

## 4.1 Proactive Capacity Planning

Detailed traffic forecasting at the hub site level can be used to proactively build capacity with a fairly high degree of accuracy. The behavior of traffic growth varies significantly between different geographies and demographics. Understanding these variations at a granular level is key to being able to accurately predict growth and even acceleration of growth due to an event like COVID-19. This is not to say that traffic forecasts related to specific one-time events can be predicted, but rather that the propensity of a site to be affected by a major event can be better understood. Higher growth sites, for instance, will likely react the most dramatically to specific events that drive sudden increases in traffic.

Various traffic forecasting and modelling approaches such as the Exponential Smoothing (ETS) methodology can be used to account for historical trends and seasonal behavior per site (Figure 4). Actual traffic observed is collected typically in monthly intervals. After some pre-conditioning of the data, it is broken down into seasonal and underlying trend components. Assembled series are then created and ETS forecasts are generated for each series. The ETS forecasted series will fall within some range of values. The mean value is used for the base traffic forecast.



**Figure 4 – Traffic Forecasting Using ETS Methodology** (Source: Comcast Capacity Management)

Once the base traffic forecast is understood, traffic impacts from future product offerings and predicted disruptors can then be incorporated to further enhance the forecast and determine capacity needs for a hub site looking forward some period of time. A twelve to eighteen-month forecast is not unreasonable if the quality of trending data is sufficient and predicted product impacts are somewhat accurate.


Not surprisingly, most hub sites which incurred the largest traffic increases during COVID-19 were already identified as higher growth and therefore had plans in place for increasing capacity, albeit a bit further into the future. The COVID-19 event acted like an accelerator of predicted traffic growth. Most of the higher growth sites were in more urbanized areas and areas with larger populations of teleworkers. More rural sites, with lower forecasted growth, tended to not be as impacted.

The proactive deployment of capacity is a preventive measure that can be used to mitigate the effects of a rapid and significant increase in network traffic. It is most effective when used in conjunction with a predictive hub site forecast.

One fairly common aspect of a proactive capacity deployment program is the use of utilization thresholds. Utilization thresholds are used to initiate capacity work before a CMTS SG becomes highly utilized. If set appropriately, this threshold would allow enough time to complete the required capacity work and provide ample headroom to account for unforeseen circumstances, such as work delays or a sudden spike in traffic growth. A more sophisticated use of utilization thresholds involves leveraging a hub site's traffic forecast to inform growth trajectories. Higher growth sites likely need to set a lower utilization threshold for capacity deployment. Thresholds could be intelligently set per hub site based on the predicted growth rate.

Proactive capacity deployment also considers the implementation of capacity solutions which provide some minimum time horizon before having to be addressed again. This varies depending on hub site growth rates. To meet a certain desired time horizon, higher growth sites require a more aggressive capacity solution. For example, a CMTS SG may breach 85% utilization and get resolved with a basic node segmentation. If the annualized growth rate of the CMTS SG is roughly 25%, it will return to the 85% utilization level again within approximately three years. If the desire is to have a longer time horizon for this CMTS SG, then a more aggressive capacity solution would be required.

Capacity solutions with longer time horizons provide proactive capacity. Deep fiber architectures and mid/high split plant upgrades are examples of proactive capacity solutions which ensure longer time horizons for a given growth rate (Figure 5). Proactive solutions tend to cost more up front but are typically less expensive over time, especially in higher growth sites where multiple node segmentations could otherwise be required every few years, and associated plant construction costs will likely continue to increase with increasing labor rates. Net present value calculations evaluating different capacity solutions have proven this in general. Another factor is customer impact. While massive plant rebuilds or deep fiber architecture solutions can sound daunting, customer impact due to capacity can be contained to one or two moments in time and then not again for several years. More short-term solutions like node segmentations would require impact to customers every few years especially in higher growth sites.



**Figure 5 – Comparison of Capacity Solutions** (Source: Comcast Capacity Management)

Hub sites built with deep fiber architectures were found to be very resilient to the rapid increases in traffic during COVID-19. Especially where N+0 amplifier cascades are implemented, node sizes for deep fiber architectures tend to be relatively small. This enables less subscribers to share the same or more capacity available on a normal node resulting in lower utilization. The percent of increase in traffic on deep fiber nodes was still relatively significant, but the impact to utilization was generally not a concern given the low utilization prior to the increase. For example, a 50% utilized node that incurred a 30% increase in traffic became 65% utilized - far below the 85% threshold of being considered highly utilized. By contrast, a 70% node that incurred the same 30% increase in traffic became 91% utilized - above the 85% threshold before being considered highly utilized. Of course, deep fiber nodes will continue to grow naturally with growth in traffic over time and will eventually need to be addressed, but that time horizon is many years down the road.

## 4.2 New and Modified Reporting

A successful capacity management program is not only proactive by nature but also very adaptable to sudden change. Adaption is largely supported by exposing the right metrics and data in order to focus and prioritize capacity work efforts in the most productive ways possible. Specialized reporting and dashboards can be very helpful to this end and do not have to be overly sophisticated. Data is the foundation, and how it is manipulated and used to tell a clear and compelling story is the key. Ideally, there are already reports and dashboards in place to help manage capacity, which can be modified during events like COVID to help adapt activities as needed.

Capacity scorecards or ranking reports are common reporting methods used to manage capacity for most service providers on a normal basis (Figure 6). These scorecards or ranking reports typically contain specific metrics and thresholds by which network capacity health is evaluated and capacity work is managed. Highly utilized nodes or CMTS SGs, for example, would be exposed and appropriate capacity mitigation work like node segmentations could be assigned and managed.

|        | Month 3            |                         | Month 2            |                         | Month 1            |                         | Legend - % of CMTS SGs |
|--------|--------------------|-------------------------|--------------------|-------------------------|--------------------|-------------------------|------------------------|
|        | %>=85% Utilization | Count >=85% Utilization | %>=85% Utilization | Count >=85% Utilization | %>=85% Utilization | Count >=85% Utilization |                        |
| Site A | 0.02%              | 2                       | 0.04%              | 5                       | 0.17%              | 22                      | 0.00%                  |
| Site B | 0.21%              | 33                      | 0.43%              | 67                      | 1.65%              | 251                     | <.05%                  |
| Site C | 0.09%              | 13                      | 0.06%              | 9                       | 0.33%              | 49                      | <.1%                   |
| Site D | 0.00%              | 0                       | 0.02%              | 2                       | 0.10%              | 11                      | <.5%                   |
| Site E | 0.12%              | 8                       | 0.16%              | 11                      | 0.78%              | 52                      | <1%                    |
| Site F | 0.03%              | 2                       | 0.17%              | 13                      | 0.34%              | 27                      | >1%                    |


**Figure 6 – Example CMTS Scorecard** (Source: Comcast Capacity Management)

During times of significant traffic growth and increases in highly utilized CMTS SGs, reporting can easily be modified or developed to provide better focus and prioritization. For example, a scorecard may be measuring how many CMTS SGs exceed an 85% utilization threshold normally. The expectation is that this number is relatively small and manageable. A sudden increase in traffic, similar to that seen during the COVID-19 event, could cause a seemingly insurmountable of CMTS SGs to spike above that 85% threshold. While efforts to increase capacity would likely happen, there are unfortunately physical limitations to how much work can be executed in a short period of time without causing other negative impacts to the network. Temporary modifications to existing scorecards could be implemented to

prioritize work efforts on all CMTS SGs exceeding 90% utilization, for example. Once the elements exceeding 90% are resolved, then focus can be placed back on those exceeding 85%. This approach was successfully adopted and implemented by Comcast to improve prioritization of work.

Other modifications to existing reporting could involve the frequency of data evaluation. Data may normally be evaluated on a weekly or even monthly cadence in the interest of removing any anomalies, such as temporary traffic spikes and CMTS SG utilization that could occur in a very specific timeframe, but subsequently decline back to pre-existing levels. Unfortunately, sudden and less temporary increases in traffic from events like COVID-19 would not be exposed initially using weekly or monthly data. Additionally, granular changes in traffic and utilization would not be known until well after the fact. A switch to more frequent reporting using daily, hourly, or even shorter time intervals is required to gain visibility to rapid increases in traffic and utilization and allow capacity management efforts to start mobilizing to the right areas faster. Furthermore, a daily review of the capacity data is warranted to ensure work teams are aligned and focused on the core priorities.

New reporting may also be warranted to drive focus and priority of capacity work efforts. At Comcast, one particular new report, called the “Minutes over Threshold (MoT)” focused on CMTS SGs and nodes with consistently high utilization over long periods of time. Some CMTS SGs incurred a level of traffic growth resulting in saturation or “flatlining,” which rarely happens normally (Figure 7). A new report was specifically created to bring even more focus and attention to these specific SGs and prioritize work accordingly. The new report not only evaluated the SG utilization level but also how many consecutive minutes those SGs remained above a certain utilization threshold. The number of minutes over threshold (MoT) metric helps identify the most severely congested SGs. Standard operating procedure would be to focus the most resources on these SGs first, then focus on SGs with lower MoT and/or lower utilization.



**Figure 7 – Example CMTS SG Saturation or “Flatlining”** (Source: Comcast Service Performance Database)

### 4.3 Capacity Acceleration

Once CMTS SG prioritization is better determined, efforts can quickly be mobilized to accelerate capacity solutions and deployments. Fast-tracking development of new innovations is very typical when new problems or challenges arise and COVID-19 was no exception for Comcast. Modifications to




existing processes and tools, accelerated bulk purchasing of materials, and rapid deployment are some other areas where Comcast focused to quickly react to the rapid increase in traffic during the COVID-19 event.

To specifically address upstream capacity challenges caused by significant teleworking and remote schooling, Comcast became highly focused on building innovative solutions to provide more upstream capacity in a condensed development timeframe. Fast-tracking development of new innovations was an essential part of our strategy during COVID-19. Resources for development and testing of new solutions were quickly re-prioritized in order to maximize speed to market of new solutions needed to expand upstream capacity. These solutions involved adding upstream channels into already fairly limited spaces in the DOCSIS radio frequency (RF) return path spectrum and implementing PMA/Octave for improved efficiency of DOCSIS 3.0 upstream performance.

Upstream channel capacity is always a challenge for cable systems limited to a 5-42 MHz DOCSIS RF return path. The typical 5-42 MHz DOCSIS return path is built with four contiguous channels at a 6.4 MHz channel-width and 64 QAM modulation roughly in the 14-40 MHz range. This is typically the maximum configuration utilized to date in most systems and enables approximately 123 Mbps of total upstream capacity. After removing all time and physical layer overhead, the usable MAC layer data rate is approximately 86 Mbps.

In order to meet the new demands on upstream capacity, qualification and testing work for the addition of a fifth and sixth upstream was ramped up. These channels would be placed in spaces above and below the four existing DOCSIS upstream channels as able depending on the channel lineup within a given cable system (Figure 8). A fifth upstream channel with 1.6 or 3.2 MHz channel-width was successfully deployed to most locations at the upper end of the 5-42 MHz range yielding a gain of 5-10% of upstream capacity, depending on channel-width. CMTS SGs at 85% utilization were effectively reduced to approximately 77-81% utilization with the addition of a fifth upstream channel. While this does not seem significant, it essentially “buys time” to complete a more impactful but more difficult-to-implement solution like a node split while striving to maintain good customer experience. A sixth upstream was also tested and deployed on a limited basis. Residing below the four existing DOCSIS upstream channels, a sixth upstream yields another 5-10% of upstream capacity but requires some changes to current DOCSIS RF return path channel lineups in most systems as well as implementation of Octave for DOCSIS 3.0 upstream.

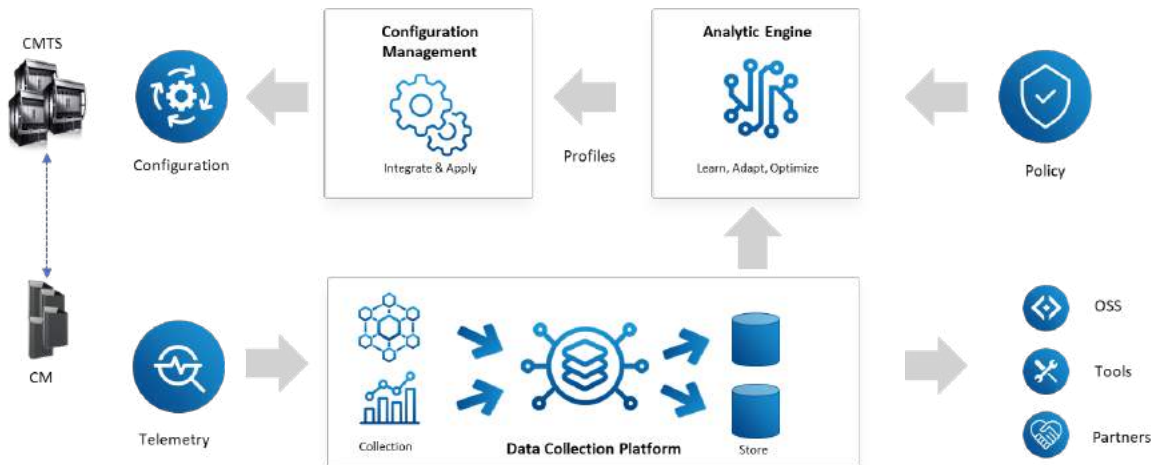


**Figure 8 – Example DOCSIS Upstream Channel Lineup With Fifth and Sixth Upstream Channel** (Source: Comcast Next Generation Access Engineering)

Regarding Octave, this was also an effort requiring accelerated development in order to provide additional upstream capacity. Originally built to support DOCSIS 3.1 downstream, Octave is an artificial intelligence (AI) platform which leverages data collection and analytics to optimize capacity based on the performance characteristics of a cable plant (Figure 9). Channel modulation can be dynamically adjusted up to a maximum of 4,096-QAM for downstream based on changing plant conditions for the DOCSIS 3.1



downstream. Without Octave, a least common denominator approach was in place typically using no more than 256 or 1,024-QAM modulation for all CMTS SGs thereby limiting potential capacity gains. Octave enabled an average gain of approximately 36% for downstream capacity. With updates to the Octave policy, we expect to be able to increase that capacity another 8% in the future.



**Figure 9 – Octave High Level System Diagram** (Source: Comcast Next Generation Access Engineering)

Long term, the plan was to always modify Octave to be used to enable more capacity for DOCSIS 3.0 upstream. COVID-19 was the impetus to accelerate the required development work and testing. At a high level, Octave operates the same way between DOCSIS 3.1 downstream and DOCSIS 3.0 upstream in that it leverages AI and automation to optimize capacity throughput given cable plant performance. The big difference though is with regard to how capacity is optimized. For DOCSIS 3.0 upstream, capacity is gained by improving mini-slot efficiency. Mini-slots are units or blocks of time used in DOCSIS upstream communication. The goal of the CMTS is to make every mini-slot productive. Depending on how various upstream parameters are configured, there may be wasted mini-slots and therefore sub-optimized data transmission. Octave was successfully leveraged for DOCSIS 3.0 upstream to evaluate the upstream parameters and optimize mini-slot utilization based on plant performance. More mini-slots are then made available for more data transmission and therefore more effective capacity. Coupled with the addition of the fifth and sixth DOCSIS upstream channels where able, Octave enabled an average gain of approximately 20%. The MAC layer data rate for the baseline four upstream channel configuration increased from 86 Mbps to 103 Mbps consistently for over 98% of CMTS SGs.


Development efforts were also accelerated on Comcast’s capacity management platform known as Enterprise Capacity and Facilities (ECAF). To facilitate the need to rapidly plan and execute capacity mitigation work, modifications were developed and implemented to add new functionality to ECAF in a short amount of time. Capacity augmentation plans, for example, would normally be entered into ECAF one at a time. To help accelerate the process, bulk entries were enabled, thereby reducing the time it took for capacity engineers to submit capacity mitigation plans. Time spent on associated documentation and administrative steps was also reduced to accommodate the much larger volumes. Daily reporting functionality was added to enable more timely and granular tracking of highly utilized SGs. This allowed teams to more quickly react to hot spots and evaluate post-mitigation effects. Automated creation of capacity mitigation plans was an existing function that was simply adjusted to a higher threshold to help manage the volume of active plans and focus on the more highly utilized SGs.

Node split planning and implementation is a key element of most capacity management programs. Depending on the specific type of split required, there are opportunities to fast-track work when circumstances require it. More simple node decombs are low-cost and easy-to-execute with little labor required and little (if any) infrastructure work. Adding optics to existing segmentable node housings is another relatively low-cost effort with a reasonably low amount of labor that can be executed fairly quickly. Construction-based node splits are relatively high-cost and labor intensive and typically require more time. Nonetheless, resources can be shifted away from other priorities, like new plant builds and metro Ethernet fiber construction, for example, and focused strictly on construction-based node splits. Node split construction permit approvals can be fast-tracked with municipalities in the interest of supporting good services to communities. Comcast effectively leveraged most of these solutions to increase node split production by 100% from normal business as usual (BAU) levels during the COVID-19 event.

Like node splits, DOCSIS channel adds are also a key part of capacity management normally. This has typically been supported over time by implementing better compression on both broadcast and narrowcast video, and then migrating available spectrum to DOCSIS. Once spectrum is available, it is fairly easy to implement on the CMTS platforms at a reasonable cost. Unfortunately, there may not be a viable means to recover further spectrum to support DOCSIS capacity near term. There may be longer term dependencies, such as conversion of video from QAM to IP-delivery. For the upstream signal path specifically, there may not be an ability to activate more DOCSIS spectrum without a costly plant upgrade. Despite the challenges, DOCSIS channel adds can easily be fast-tracked when able to make a practically instant impact on capacity utilization. As mentioned previously, adding a fifth upstream adds approximately 5-10% of capacity depending on channel-width and modulation configuration. This effectively results in a reduction of utilization of ~4-9%. Adding eight 6 MHz downstream channels to an existing 32-channel SG reduces utilization by 18-23% on average.

Regardless of the type, most capacity work requires some amount of hardware and software investment. This is where solid communications and relationships with a company's finance and procurement organizations come into play. Vendor relationships are also a crucial part of the overall supply chain. Normal planning should be in place to set internal budget forecasts and provide some degree of material forecasting and revenue planning for vendors. Rapid increases in node splits and DOCSIS channel additions drive the need for more hardware and DOCSIS licensing sooner. Capital was efficiently accelerated to support higher volumes of work needed. Vendors mobilized resources to support the higher demand even amidst challenging workplace restrictions. Not surprisingly, acceleration of hardware and software is most successful when a proactive site-level plan is already in place. The effort is mostly a matter of accelerating existing plans and making slight adjustments where needed.

The numerous capacity solutions and efforts described here have proven successful in mitigating highly utilized CMTS SGs. From its peak in mid April, upstream SG utilization was reduced approximately 70% by early July. Likewise, downstream SG utilization was reduced approximately 90% by early July and lower than pre-COVID-19 levels tracked in February (Figure 10).



**Figure 10 – Highly Utilized CMTS SG Trending** (Source: Comcast Capacity Management)

## 5. Conclusion

COVID-19 is inarguably one of the most impactful and disruptive events in recent world history. Much of the world's population has been forced to react and change lifestyles in an effort to mitigate the spread and effects of the virus. Many of these lifestyle changes required people to work and school from home at a scale never before observed. This resulted in massive increases in network traffic on service provider networks in many areas. For MSOs, this increased traffic manifested itself as higher and in some cases much higher utilization on CMTS SGs potentially causing impact to customer service and experience.

Much was learned from the COVID-19 event in terms of how to successfully manage the access network and ensure good customer service. Above all, a solid capacity management program with the right people, processes, and tools is fundamental at all times and not just when a pandemic hits. It is too late to put such a program in place when that happens.

A good capacity management program should have a bias toward proactivity. It should strive to maintain some amount of available capacity at all times to account for sudden increased demand on the network. There is simply never enough time to react to sudden events that impact network utilization. This necessitates the need implement proactive capacity efforts, particularly around predictive traffic forecasting and capacity planning. This is not suggesting that timing for events like COVID-19 can be predicted. Rather, understanding the growth potential of nodes and CMTS SGs in different geographic areas based on historical behavior allows capacity management teams to appropriately plan some amount of proactive capacity with the ability to adjust quickly if needed.

Proactive deployment of capacity naturally follows from predictive traffic forecasting and planning. Normally high growth hubsites are likely to be the most impacted during an event that drives a rapid increase in utilization. Capacity plans for hub sites need to take into account this potential for growth.

More aggressive and proactive capacity solutions like deep fiber or plant upgrades should be deployed into sites with higher growth. Lower growth sites may be fine with less aggressive solutions. In most cases, proactively deployed capacity will be sufficient to handle extreme events that drive rapid increases in network utilization. In others, additional reactive measures may be required. Either way, there is no chance for success without proactive capacity in some form.

As an event unfolds driving significant network utilization growth, capacity management programs need to implement a variety of reactive measures building from a foundation people, processes, and tools that should already be in place. Reactive measures start with enhancements in reporting such as scorecards and dashboards to better identify problems and prioritize capacity mitigation work. Altogether new reports may need to be developed using existing or new data feeds. Development resources may need to be shifted to focus on building new functionality in reporting and tools.

Once impacts from sudden increases in network utilization are better understood leveraging enhanced reporting and data, efforts can be made to accelerate and prioritize capacity work. Fast-tracking development of new innovations such as Octave are effective ways to rapidly deploy additional upstream capacity especially to areas with high utilization without any additional hardware or physical work required. New functionality in capacity management platforms like ECAF enables higher velocity of capacity solution execution.

Acceleration of existing capacity plans like node splits and DOCSIS channel adds are extremely effective and generally easy to execute depending on the specific circumstances for a given node, CMTS SG, or hub site. Node split volumes were doubled at Comcast as part of efforts to accelerate capacity. Additional or accelerated purchasing of hardware and software may be needed.

Solid partnership with finance, purchasing, and vendor teams is needed here. Rapid execution of physical work in support of capacity solutions requires commitment and dedication from construction, operations, and government affairs teams to name a few. Ultimately, victory in the battle against rapid growth in network utilization happens when all teams and organizations come together on the aligned goal to accelerate capacity,

Hopefully, we will not see another pandemic like COVID-19 for a very long time. But, the reality of the matter is that unforeseen events of varying magnitude and impact will always happen. That's life. For those events that drive significant impact on service provider networks even if only on a limited scale, the lessons learned from the COVID-19 about capacity management will always be valuable and applicable.

# Acknowledgements

I would like to thank all of the Comcast capacity management, engineering, and operationsteams and partners for their hard work, effort, and relentless passion. Their perseverance was key to successfully managing capacity during extreme circumstances.

I would also like to acknowledge the great work done by Dan Rice, Vice President of Access Engineering at Comcast Cable, and the extended Octave team for their efforts in rapidly adjusting focus to develop and test new upstream capacity solutions for accelerated deployment in response to extreme network growth incurred after lock-downs imposed at onset of COVID-19 . Dan's team enabled the accelerated deployment of fifth and sixth DOCSIS upstream channels in the 5-42 MHz DOCSIS return path. Additionally, Dan's team was able to successfully modify and launch PMA/Octave for DOCSIS 3.0 upstream capacity optimization.

# Abbreviations

|        |                                                |
|--------|------------------------------------------------|
| AI     | Artificial Intelligence                        |
| BAU    | Business As Usual                              |
| CM     | Cable Modem                                    |
| CMTS   | Cable Modem Termination System                 |
| DOCSIS | Data Over Cable System Interface Specification |
| ECAF   | Enterprise Capacity and Facilities             |
| ETS    | Exponential Smoothing                          |
| Mbps   | Megabits per second                            |
| MHz    | MegaHertz                                      |
| MoT    | Minutes over Threshold                         |
| MSO    | Multiple-System Operator                       |
| OFDM   | Orthogonal Frequency Division Multiplexing     |
| OSS    | Operating Support System                       |
| OTT    | Over The Top                                   |
| PMA    | Profile Management Application                 |
| QAM    | Quadrature Amplitude Modulation                |
| RF     | Radio Frequency                                |
| SG     | Service Group                                  |
| VOD    | Video On Demand                                |
| VoIP   | Voice over IP                                  |
| VPN    | Virtual Private Network                        |

# **Small Cell Traffic Engineering**

## **How Many Small Cells are Needed for Proper Coverage?**

A Technical Paper prepared for SCTE•ISBE by

**John T Chapman**  
CTO Cable Access, Fellow  
Cisco Systems  
300 East Tasman Drive, San Jose, CA  
408-526-7651  
jchapman@cisco.com

# Table of Contents

| Title                                                | Page Number |
|------------------------------------------------------|-------------|
| 1. Introduction.....                                 | 4           |
| 2. Small Cells per MacroCell.....                    | 5           |
| 2.1.    CBRS Quick Primer .....                      | 5           |
| 2.2.    Small Cell and Macrocell Coverage.....       | 6           |
| 2.3.    Mathematical Model I.....                    | 8           |
| 2.4.    Mathematical Model II.....                   | 10          |
| 2.5.    Small Cell Radios per Macrocell Example..... | 11          |
| 2.6.    Cost Implications.....                       | 11          |
| 3. Small Cells per HFC Fiber Node.....               | 13          |
| 3.1.    DOCSIS-Attached Small Cells.....             | 13          |
| 3.2.    Real World Case Study.....                   | 14          |
| 3.3.    Mathematical Model.....                      | 15          |
| 3.4.    Small Cells per Fiber Node Example.....      | 17          |
| 3.5.    Kindred Spirits.....                         | 18          |
| 4. Is DOCSIS Up to the Job? .....                    | 19          |
| 4.1.    CBRS Use Case for HFC.....                   | 19          |
| 4.2.    The Future Potential of DOCSIS Capacity..... | 19          |
| 4.3.    SYNC – 1588 and SyncE Support.....           | 21          |
| 4.4.    LLX – Low Latency Xhaul.....                 | 21          |
| 5. Conclusion.....                                   | 23          |
| Abbreviations .....                                  | 24          |
| Bibliography & References.....                       | 25          |

## List of Figures

| Title                                                        | Page Number |
|--------------------------------------------------------------|-------------|
| Figure 1 – Small Cell Deployment Model.....                  | 4           |
| Figure 2 – Macrocell versus Small Cell Sites.....            | 6           |
| Figure 3 – RF Cell Radius Comparison .....                   | 7           |
| Figure 4 – Model of Small Cell per Macrocell (Max).....      | 8           |
| Figure 5 – Coverage of a Small Cell with Ideal Overlap ..... | 9           |
| Figure 6 – Model of Small Cell per Macrocell (Min).....      | 10          |
| Figure 7 – MVNO Roaming Costs .....                          | 12          |
| Figure 8 – Strand-Mount Small Cell Deployment.....           | 13          |
| Figure 9 – Real World Analysis of xHaul .....                | 14          |
| Figure 10 – Model of Small Cell per HFC Optical Node .....   | 15          |
| Figure 11 – DOCSIS Plans for Spectrum Increases.....         | 20          |
| Figure 12 – SYNC Operating Model .....                       | 21          |
| Figure 13 – LLX Operating Model.....                         | 21          |
| Figure 14 – Macrocells and Small Cells in Deployment .....   | 23          |



## List of Tables

| <b>Title</b>                                               | <b>Page Number</b> |
|------------------------------------------------------------|--------------------|
| Table 1 – RF Cell Radius Comparison .....                  | 7                  |
| Table 2 – RF Cell Radius Comparison .....                  | 11                 |
| Table 3 – Real World Case Study of xHaul over DOCSIS ..... | 14                 |
| Table 4 – Small Cells per Fiber Node .....                 | 17                 |

# 1. Introduction


The world as we know it in cable is about to change.

Today's cable operators are tomorrow's mobile operators

This is being driven by a multitude of factors. For one, it is a way of significantly increasing revenues. All the cable broadband customers have cell phones and are often paying two different service providers, one for wireline and one for wireless. Does that really make sense?

Remember the telephone network? That country wide collection of twisted pair copper and T1 lines? It does not exist anymore. Telephony is an application that runs over an IP network. There is no Wi-Fi network. Wi-Fi is a radio gateway that exists at the end of an IP network that provides end point connectivity.

Today, we still have separate and distinct mobile and cable networks. But not for long. These too will soon also become applications running over an IP network. As mobile moves to 5G and deploys small cells running at higher frequencies, that small cell will become a Radio Frequency (RF) gateway that converts between wireless and a wireline IP network with traffic tunneled to a 5G core (5GC).



**Figure 1 – Small Cell Deployment Model**

The same is true with cable. DAA (Distributed Access Architecture) [1][2] nodes are simple RF gateways that convert between digital fiber and RF over coax. This results in fiber to the neighborhood and DOCSIS to the door. With DAA, there is IP over Ethernet over fiber, followed by IP over Ethernet over DOCSIS over coax, followed by IP over Wi-Fi or Ethernet.

This leads us to another important premise.

Behind every great wireless network is a great wireline network

Every radio that creates its part of the mobile wireless network must be connected to a wireline network. In Figure 1, we see the deployment model published by the city of Danville, CA [1]. Notice that the small cell coverage is a fraction (0.2% to 2%) of the macrocell coverage. That is a big difference.

How many small cells will it take to provide the equivalent coverage of a macrocell? How will those get connected? What is the impact to cost? Those are the questions we will tackle in this white paper.

## **2. Small Cells per MacroCell**

### **2.1. CBRS Quick Primer**

Going forward, a very important frequency band of interest to the cable and mobile industry is the Citizens Broadband Radio Service (CBRS) which is a 150 MHz band located between 3550 MHz and 3700 MHz. This RF region is also known as band 48.

The CBRS spectrum is allocated amongst three tiers. The first tier and highest priority are the incumbents like the US Navy and they get first use. The second tier is a licensed spectrum known Priority Access License (PAL) and is composed of operators who have bought local spectrum in the lower 100 MHz through a system known as the Spectrum Access System (SAS). The third and lowest priority tier is unlicensed (which means it is free) known as General Authorized Access (GAA). GAA is composed of 80 to 100 MHz of spectrum that is dynamically allocated out of the 150 MHz CBRS band. This is where “Private LTE” will exist.

CBRS spectrum tends to get allocated out in sub-bands of 10 MHz. Sub-bands can be combined together. As an example of performance, an example outdoor small cell radio would have 40 MHz max RF bandwidth, 200 Mbps max throughput with one-watt Effective Isotropic Radiated Power (EIRP). That works out to 5 bits/Hz. Note how well 100 Mbps to 200 Mbps throughput matches to a DOCSIS CM.

## 2.2. Small Cell and Macrocell Coverage



**Macrocell on a Tower**



**Small cell on a Lamp Post**

**Figure 2 – Macrocell versus Small Cell Sites**


Radio coverage depends on three fundamental factors:

- Frequency
- Power
- Height and physical interference

The higher the frequency, the higher the propagation loss and hence the smaller the coverage area. Higher frequencies generally tend to experience higher loss from objects in the path of propagation. The higher the power, the longer the propagation. And the higher the antenna, the less likely it will have physical interference from tree, hills, and walls.

Macrocells tend to be lower in frequency, higher in power, and mounted at higher heights. Small cells tend to be higher in frequency, lower in power, and mounted at lower heights. These combined factors drastically lessen the coverage ratio of a small cell compared to a macrocell. Some typical installations are shown in Figure 2.

There are other factors to be sure such as modulation and error-correction that impact channel capacity and toleration to distance, but they are more secondary factors.



**Figure 3 – RF Cell Radius Comparison**

Figure 3 shows the ratio of a macrocell radius to a small cell radius [4]. The cell radius at 700 MHz is normalized to one. The frequencies at or adjacent to 700 MHz are used for Long Term Evolution (LTE) mobile systems. The values from Figure 3 are also shown in Table 1.

**Table 1 – RF Cell Radius Comparison**


| Band    | Service | Cell Type  | Relative Radius |
|---------|---------|------------|-----------------|
| 700 MHz | LTE     | macrocell  | 1.0             |
| 3.5 GHz | CBRS    | small cell | 0.09            |
| 28 GHz  | mmWave  | small cell | 0.003           |

MoffettNathanson states that the numbers are just illustrative; they assumed the same power level and antenna height for each spectrum band. When compared to Figure 1, where the ratio was 0.2% to 2%, the numbers in Table 1 which are based on free air may actually be optimistic. This means in practice, with lower heights and lower powers, the relative radius may be lower and the calculated small cell count will be higher.

Conversely, small cells can take advantage of the latest technologies including beam forming and MIMO to extend their reach. It also makes a big difference whether the small cells are deployed on a stand-mount in a crowded downtown area or in a rural area or a rooftop or hilltop. Still, it is important to start somewhere to establish a baseline, and then that baseline can be adjusted up or down.

With these data points and some high school algebra, we can calculate how many small cells it would take for an equivalent macrocell footprint.

### 2.3. Mathematical Model I



**Figure 4 – Model of Small Cell per Macrocell (Max)**

The following is a mathematical model of how many small cells it would take to achieve 100% coverage equivalence to a macrocell. The analysis assumes free air, so real-world results will vary.

The models in this white paper are derived here and do not reply on previous literature.

Let's first assume that a cellular radio cell RF coverage is a perfect circle. It is not enough to divide the area of the small circle of the small cell into the large circle of a macrocell since this would assume no overlap. Circles that have no overlap also have gaps in coverage as shown in Figure 4. Instead, if we assumed a maximized square inside the small circle, then the squares can then be tiled to create 100% coverage.

In practice, there is almost always either too much overlap of cell radius, resulting in some loss of coverage efficiency, or not enough overlap, resulting in gaps of coverage. Also, the macrocell itself is overlapping or gapping with other neighboring macrocells. None the less, this model provides a starting point where the optimized configuration which can then be adjusted up or down, based on deployment considerations.

The following should bring back pleasant memories of high school algebra. If those memories are more of a nightmare, then just skip this section. I won't be offended.

Let's define:


$R$  = Radius of the macrocell

$r$  = radius of the small cell

The coverage of the macrocell area, based on a large circle, is:

$$\text{macrocell coverage} = \pi R^2 \quad \text{formula (1)}$$

That was easy. Now the next part. The coverage of the small cell area is based on a maximum sided square inside the small circle. The radius of the small circle is equal to the distance from the center of the square to the corner of the square. That radius is part of a small triangle in the corner of the square. This is shown in Figure 5.



**Figure 5 – Coverage of a Small Cell with Ideal Overlap**

The Pythagorean Theorem provides the length of the triangle side.

$$a^2 + b^2 = r^2$$

If  $b = a$ , then

$$a^2 + a^2 = r^2$$

$$2a^2 = r^2$$

$$a^2 = r^2/2$$

$$a = r/\sqrt{2}$$

The length of the side of the square is twice this value.

$$\text{box side} = 2a = 2r/\sqrt{2} = \sqrt{2}r$$

The area of the square, and hence the small cell coverage we are looking for, is the square of the box side.

$$\text{small cell coverage} = (2a)^2 = (\sqrt{2}r)^2$$

$$\text{small cell coverage} = 2r^2 \quad \text{formula (2)}$$

We can now compute the ratio of macrocell coverage, formula (1), to small cell coverage, formula (2).

$$\#SC \text{ per } MC = \text{macrocell coverage} / \text{small cell coverage}$$


$$\#SC \text{ per } MC = \pi R^2 / 2r^2$$

$$\#SC \text{ per } MC = \pi/2 (R/r)^2 \quad \text{formula (3)}$$

If R is normalized to 1, then we get a simplified version that looks like

$$\#SC \text{ per } MC = \pi/2r^2 \quad \text{formula (4)}$$

## 2.4. Mathematical Model II



**Figure 6 – Model of Small Cell per Macrocell (Min)**

In the previous section, we placed overlapping small cells into a macrocell radius. To be a bit more optimistic, what if that macrocell was also just perfectly overlapping with other macrocells. Then, the calculation would simplify to small squares filling a big square. This approach provides a range between two extreme cases – one with no overlap and one with perfect overlap.

Formula (2) provided the area of a square based on its radius. This can be applied to

$$\#SC \text{ per } MC \text{ min} = \pi R^2 / \pi r^2$$

$$\#SC \text{ per } MC \text{ min} = (R/r)^2 \quad \text{formula (5)}$$

Normalizing the larger radius R to R = 1:



$$\#SC \text{ per MC min} = 1/r^2 \quad \text{formula (6)}$$

There, that was not so bad. Now let's plug in some numbers and see what happens.

## 2.5. Small Cell Radios per Macrocell Example

Now let's take the values from Table 1 and insert them into formula (4) and formula (6) to get Table 2.

**Table 2 – RF Cell Radius Comparison**

| Band    | Service | Cell Type | Relative Radius | # Radios           |
|---------|---------|-----------|-----------------|--------------------|
| 700 MHz | LTE     | MC        | 1.0             | 1                  |
| 3.5 GHz | CBRS    | SC        | 0.09            | 125 to 200         |
| 28 GHz  | mmWave  | SC        | 0.003           | 110,000 to 175,000 |

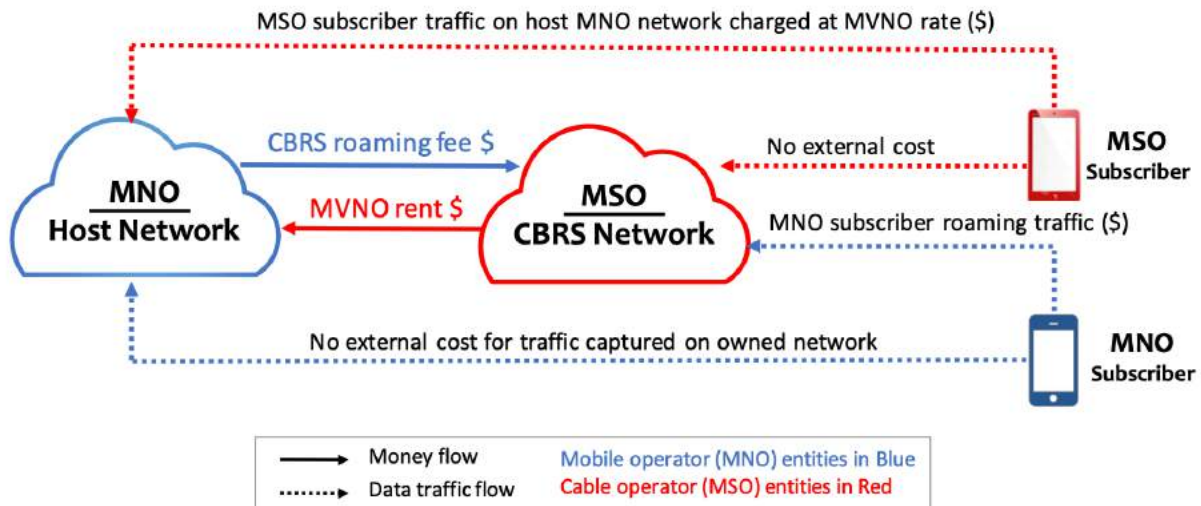
The results in Table 2 state that for the CBRS band which located at 3.5 MHz, it would take about 125 to 200 small cell radios to match the coverage of one LTE macrocell at 700 MHz. Note that the results in Table 2 have been rounded off in alignment with rules for significant figures.

28 GHz is the band of interest that Verizon is currently using for 5G mmWave [5]. It would take 110,000 to 175,000 small cell radios to replace one LTE macrocell. That's a lot of radios. Even at only 1% coverage, that would be 1750 radios.

Now, if the strategy is to not achieve 100% coverage and to also rely on the LTE network as a background network, then these numbers can be scaled down. Conversely, small cells that operate at a both a higher frequency and a lower installation height than LTE towers are also more likely to get blocked by trees or walls. Under those circumstance, the number of radios in Table 2 may actually need to be increased.

## 2.6. Cost Implications

Depending upon business needs, operators may not need a 100% replacement of the macrocell. For example, when AT&T used LTE femtocells to supplement LTE, those femtocells were only needed in areas of weak or no coverage. Conversely, as US cable operators enter into mobile virtual network operator (MVNO) arrangements, they will be financially motivated to put CBRS small cells in areas where there is any significant amount of traffic so that they do not have to pay back carriage charges to their MVNO partner that manages the LTE network. The flow of money from roaming fees for an MSO acting as an MVNO and its host network partner is shown in Figure 7 which is from the discussion in [6].



**Figure 7 – MVNO Roaming Costs**

This paper has shown that it can take 200 CBRs small cells to cover an area equivalent to the area covered by one LTE macrocell. Even if the goal was only 50% coverage, that is still 100 small cells. And if the goal was to build a new radio network that cost equal to or less than the LTE network, then the install cost of each small cell would have to be 1% of what it took to install an LTE macrocell.

So, it has to be almost free.

A typical macrocell installation cost is \$20,000 to \$50,000 to rent a tower, get permits, trench and run power and fiber, and mount a radio. If economics need to drop to 1% with 50% coverage, then the small cell deployment installation would have to be \$200 to \$500 per small cell. Is that possible? Well, that is about the cost it takes to deploy a residential CM if there is a truck roll. So, if a small cell was as cheap and easy to deploy as a CM, then yes. As a case in point, Wi-Fi is almost free to install because it is included in the CM install. So, if a small cell cost were also included in the CM cost, it could see similar cost structures.

So, it can be almost free.

There is another important consideration, and that is xhaul. Xhaul refers to backhaul, midhaul, or fronthaul. For LTE, the eNB generates GTP (GPRS Tunneling Protocol) encapsulated packets ready for backhaul. In 5G, the equivalent gNB is divided into an RU (Radio Unit), DU (Distributed Unit) and a CU (Centralized Unit). DOCSIS can support either the backhaul from a gNB (RU+DU+CU) or midhaul from a RU+DU. DOCSIS does not have the bandwidth or latency to support fronthaul from an RU.


Today's macrocells often have a fiber backhaul. If you had to deploy 100x more radios at 100 new locations, is it cost effective to run 100 more fibers? If the broadband cable team could not justify the cost of the additional fiber for residential broadband, how is it that the company business case changes such that it works for small cells?

These economics are now driving cable operators such as Charter and Telecom Argentina [7], Cogeco [8], Shaw [9], Cox [10], and Altice [11] amongst others, to deploy mobile backhaul over DOCSIS. I have also previously discussed this business case in my blog post [11].

In the next section, we will look at some real-world results between using coax direct connect versus running fiber a few blocks.

### 3. Small Cells per HFC Fiber Node

#### 3.1. DOCSIS-Attached Small Cells



**Figure 8 – Strand-Mount Small Cell Deployment**

Macrocells have traditionally been on towers and backhauled by fiber, but there are not enough towers or fiber for the new 5G small cells.

Cellular radios require the following basic requirements:

1. Site/Location
2. Power
3. Backhaul

The following requirements are also useful based upon the installation

4. Timing support
5. Low latency

DOCSIS over HFC (Hybrid Fiber Coax) supports all these requirements. HFC passes 93% of USA HHP and provides strand mount or in-home. Altice, Cox and Shaw have already deployed strand-mount small cell Xhaul over DOCSIS (>20,000).

Figure 8 from [13] shows a typical strand mount small cell. There is a directional coupler that couples power and RF from the coax plant. This goes to a strand mount CM which then connects to a strand mount small cell.

It should be noted that the HFC plant has to be designed to allow for the power drop across the directional coupler and for the additional power draw of the small cell as well as the additional power dissipation of the cable modems and small cells.

An alternative mounting configuration is to locate the small cell in the subscriber residence. This again provides location, power, and backhaul.

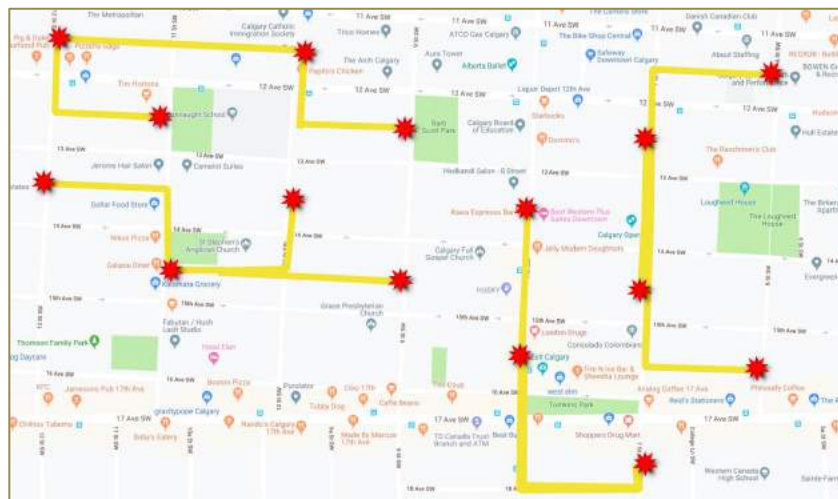
So, what does this cable plant look like, and how many small cells would a fiber node in an HFC plant be expected to support?

### 3.2. Real World Case Study

A study was done by Shaw Communications in 2019 [9] that compared a fiber backhaul solution to a coax/DOCSIS backhaul solution for a set of 15 small cells deployed across 13 node locations. In this study, only time and construction costs were considered. The deployment diagram is shown in Figure 9.

For fiber backhaul case, fiber was run (yellow lines) from the small cell location (red stars) to the node location and connected to available dark fiber. Only the cost of this short fiber run was included. The cost of the fiber from the node to the hub was not included. If it was, the difference would even be more dramatic.

For the coax backhaul case, the small cells were powered from the HFC plant with a 10' coax drop and a connected with a bidirectional coupler. The strand mount small cells did not require any permitting or have any access issues and no civil build was required.



**Figure 9 – Real World Analysis of xHaul**


The results are in Table 3 and are astonishing. The cost of installation dropped to 1% and the time-to-deploy became 20x faster. Using an available HFC plant instead of new fiber was a clear winner.

**Table 3 – Real World Case Study of xHaul over DOCSIS**

| Backhaul Option | Broadband Fibers | Construction Cost                       | Time to Build                       |
|-----------------|------------------|-----------------------------------------|-------------------------------------|
| Fiber           | 1                | \$182,500                               | 4-6 months                          |
| Coax            | 0                | \$1,500                                 | 1 week                              |
| <b>Savings</b>  |                  | <b>1% of the cost<br/>99% reduction</b> | <b>20x Faster<br/>95% reduction</b> |

Now that we have established a deployment and business case baseline, let's figure out how many small cells might fit into an HFC plant.

### 3.3. Mathematical Model



**Figure 10 – Model of Small Cell per HFC Optical Node**

So how many small cells would it take to cover an area that is currently covered by a DOCSIS service group that is associated with one fiber node?

That depends on the size of the plant. HFC plants are described by their amplifier depth. An N+0 plant is a node plus zero amplifiers. An N+5 plant is a plant with a node followed by five consecutive amplifiers. Note that the signal path in an N+5 plant is constantly being split, which creates the area coverage. Thus a single node could have as many as 25 to 40 total amplifiers across all paths.

Let's define the following variables:

S = the average span distance between node-to-amp and amp-to-amp

M = the number of amplifiers in the longest cascade, as in "N + M"

R = The resulting node radius

A diagram of an N+2 plant is shown in Figure 10. A radius line "R" is drawn from the node to the end of the last cable segment. Note that a two-amplifier system has three cable segments. Coax is usually run down streets which leads to rectangular shapes. It can be seen that the horizontal segments fully line up with the "x" axis and the vertical segments fully line up with the "y" axis. These observation leads to the following formula:

$$x + y = (M + 1)S$$

So, the total horizontal "x" and vertical distance "y" is equal to the average segment length "S", times the number of segments which is one more than the number of amplifiers "M". If we assume that the triangle is an isosceles triangle where  $x = y$ , then:

$$x + x = (M + 1)S$$

$$2x = (M + 1)S$$

$$x = (M + 1)S/2 \quad \text{formula (7)}$$

Checking back in with Pythagoras, we also know from Figure 10 that

$$x^2 + y^2 = R^2$$

Since  $y = x$

$$x^2 + x^2 = R^2$$

$$2x^2 = R^2$$

$$x^2 = R^2/2$$

$$x = R/\sqrt{2} \quad \text{formula (8)}$$

Equating formula (7) and formula (8) yields

$$R/\sqrt{2} = (M + 1)S/2$$

$$R = (M + 1)S/\sqrt{2} \quad \text{formula (9)}$$

Formula (9) calculates the radius of a fiber node serving area. An alternate form of formula (9) is

$$R = 0.707(M + 1)S \quad \text{formula (10)}$$

As a rule of thumb, that means the radius of an HFC plant is approximately equal to 70% of the total span length.

Formula (9) can be put into formula (3) to calculate the number of small cells that are needed to cover the fiber node serving area.

$$\#SC \text{ per FN} = \pi/2 (R/r)^2$$

$$\#SC \text{ per FN} = \pi/2 \left( (M + 1)S/\sqrt{2}r \right)^2$$

$$\#SC \text{ per FN} = \pi/4 (S/r)^2 (M + 1)^2 \quad \text{formula (11)}$$

Formula (11) calculates the number of radios in a fiber node serving area for an average amplifier span distance of S, and amplifier count of M, and a small cell serving radius of r.

What we can observe from this formula, is that when the small cell radius is equal to the node segment distance ( $S = r$ ) on a deep fiber plant ( $M = 0$ ), only one small cell radio is needed per node. This makes sense.

As HFC plant grows beyond  $N + 0$ , the number of radios required is proportional to the square of the number of amplifiers plus 1. If the small cell radius is larger or smaller than the span distance, then there is also an additional relationship to the square of the ratio of the span distance,  $S$ , to the small cell radius,  $r$ , as well.

### 3.4. Small Cells per Fiber Node Example

Formula (9) and formula (11) use the average span length of an HFC plant. The average span length depends on how many spans there are. So, to construct a table of results, we need to modify the average span length for each N+M plant design. Here is the methodology used to construct Table 4.

- An N+0 HFC plant can actually run the fiber node output amplifier closer to saturation and achieve longer distances. This example will add 10% length for N+0.
- An N+M HFC plant will have an initial span length. A typical value is 1000' (300m).
- Each amplifier adds to the noise floor which reduces MER, so the span length is shortened slightly for each subsequent span. This is a minor consideration as each plant is unique. The number of tap groups per span also may change. The table uses a manual entry for total plant length for each case and calculates an average.

Note that in deployment, the fiber nodes that drive an N+M plant are often called BAU (Business as Usual) nodes and the nodes that drive an N+0 plant are often called SHO (Super High Output) nodes.

To make comparisons easy, the small radius has been chosen to be 50%, 100% and 200% of the initial span length. This will illustrate the squared relationship of the span length to cell radius ratio and the number of small cells needed.

The results of this study are in Table 4. There are some interesting takeaways. First, there is not one rule for all deployments. It is not that one small cell per fiber node always works. Instead, it depends upon the size of the HFC plant. There could be anywhere from one to 20 or even 80 small cells required to cover a plant area.

Does that make sense? Well, if you had a 500 HHP (households passed) N+5 node with 50% of those households being cable-mobile customers, and the small cell was located inside the house, then that would be 250 small cells per node. So, by that argument, the results in Table 4 are conservative or should be looked at as strand-mount numbers.

Note that these calculations are for free air, which implies line-of-site with no hills, trees, or walls, so more small cells are likely to be needed in a real-world deployment.

**Table 4 – Small Cells per Fiber Node**

|        |            | Small cell radius r: |          | 500      | 1000     | 2000     |
|--------|------------|----------------------|----------|----------|----------|----------|
| M Amps | Total Span | Node Radius          | Avg Span | # radios | # radios | # radios |
| 0      | 1100       | 778                  | 1100     | 4        | 1        | 1        |
| 1      | 2000       | 1414                 | 1000     | 13       | 3        | 1        |
| 2      | 2850       | 2015                 | 950      | 26       | 6        | 2        |
| 3      | 3700       | 2616                 | 925      | 43       | 11       | 3        |
| 4      | 4600       | 3253                 | 920      | 66       | 17       | 4        |
| 5      | 5460       | 3861                 | 910      | 94       | 23       | 6        |
| 6      | 6300       | 4455                 | 900      | 125      | 31       | 8        |

The caveats are all around height of the small cell, trees, hills, and walls, as well as HFC plant design. Individual results will vary. However, the basic principle remains.

Reach out to me on LinkedIn or email if you would like a copy of these calculations in a spreadsheet.

### **3.5. Kindred Spirits**

Here is a philosophical-techno thought.

The mobile small cell RF downlink and the HFC node RF downstream both try and do the same thing, and that is to propagate RF through a media. One does so through the air and the other through coax. The one through the air runs into trees, hills, and walls, while the one on coax goes past trees, over hills, and through walls.

Both cover a serving radius. The one that propagates through air goes wherever air goes, which is everywhere. The one that propagates through coax goes only where coax goes, which is strand, underground coax and customer premises wiring.

Both have the same propagation limitations based on frequency and power, although that for a given frequency, the loss per unit distance is much higher through coax than that through air. They both use similar electronics and modulations, and sometimes similar frequencies. Where the coax distribution loses power to taps, the air distribution loses power to trees and walls.

Each has similar goals but has its own struggles and accomplishments.

So, it makes some sense, that the radius of a small cell might be on the order of a span of coax.

Note that coax express runs are one dimensional with no taps and go longer distances. They do not apply to this study or analogy.



## 4. Is DOCSIS Up to the Job?

### 4.1. CBRS Use Case for HFC

A common misconception is that if you deploy many small cell radios, the network bandwidth required is an aggregate the sum of all that bandwidth. CBRS radios can support 50 Mbps per 10 MHz of spectrum (advanced configurations may get to 100 and 200 Mbps by combining spectrum). If you take a serving area that is supported by one large radio at 100 Mbps and replace it with 50 small radios, each at 100 Mbps, would you need 50 radios x 100 Mbps = 5 Gbps of data capacity? No, you would not.

This is about capacity versus connectivity. Small cells, like CMs, need to be at a certain location like a home to provide coverage, regardless if the network is idle or active. When replacing one macrocell with say 50 small cells, it is about coverage and connectivity first and capacity second. Those 50 small cells will connect to the same end points – that would be you and your family, with your smart phones, laptops, and IP set-top boxes.

For these reasons, not all small cells will operate at peak capacity at the same time. We see that with cable modems deployments where there is a large over-subscription of offered data bandwidth with respect to the actual data bandwidth used. This also occurs in fiber installations where the backhaul interface from a fiber OLT has a typical concentration of 20:1. This is traffic engineering at work. Small cells can take care of this over-subscription as well. The inverse of over-subscription is concurrency. A 20:1 over-subscription would support 5% concurrency.

We also know that the radios peak throughput decreases with distance. That means that different users will get different peak throughputs. In traffic engineering, that is managed with an average throughput per user. A more reasonable way to calculate network loading of small cells (SC), taking into account 75% average throughput per user and a generous 20% concurrency (5:1 over-subscription), would be:

$$SC \text{ Network Loading} = \# SC * 100 \text{ Mbps peak} * \text{avg capacity} * \text{concurrency} \quad \text{formula (12)}$$

$$SC \text{ Network Loading} = 50 SC * 100 \text{ Mbps peak} * 75\% \text{ avg capacity} * 20\% \text{ concurrency}$$

$$SC \text{ Network Loading} = 750 \text{ Mbps}$$

Now, if those small cells are added to the home network with existing subscribers and existing devices, this bandwidth may already be accounted for and very little increase may be seen. If those small cells are located outdoors and pick up new subscribers and new devices, then 750 Mbps is about one-half of an OFDM channel backhaul.

As a gut check, the CMs in deployment typically provide 100 Mbps to 1 Gbps of service. If in a deployment of 200 devices (CMs), you added 50 more devices (small cells) at 100 Mbps, the difference would be incremental, not monumental.

### 4.2. The Future Potential of DOCSIS Capacity

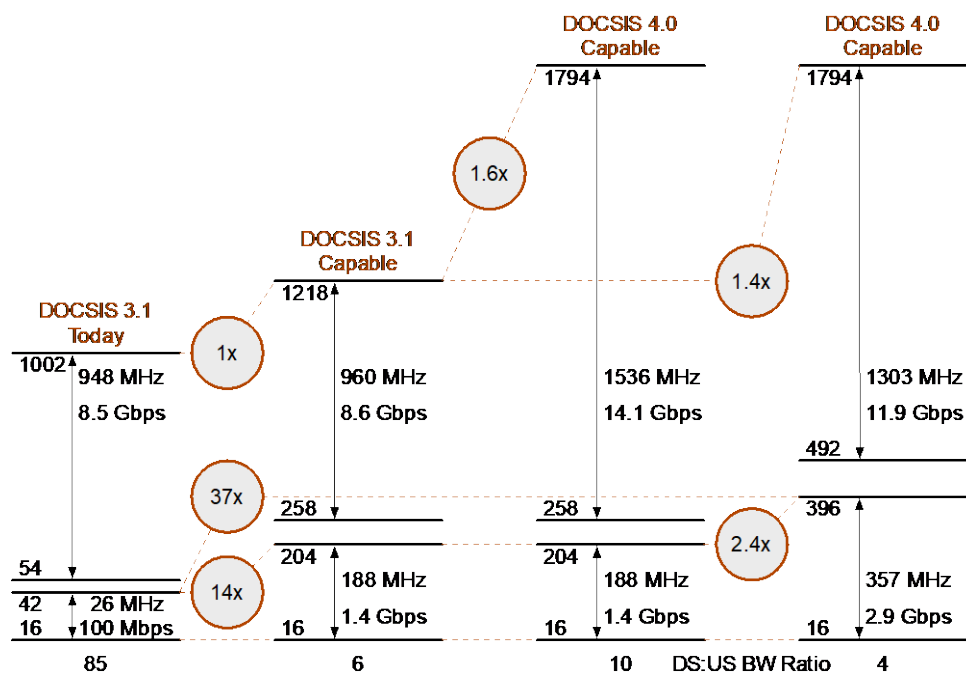
A common question often comes up. Does DOCSIS have enough bandwidth to support a small cell infrastructure? DOCSIS has two growths paths. These are spectrum and segmentation. Let's look at both.

Figure 11 (original by author) shows select downstream and upstream spectrum usage for DOCSIS 3.1 and the upcoming DOCSIS 4.0 [14]. Today the downstream spectrum is shared with legacy MPEG-TS

video, so the entire downstream is not available for DOCSIS. Actually, a good way to increase DOCSIS spectrum is through legacy video reclamation.

Today's deployed DOCSIS downstream paths are 1 to 2 Gbps in capacity. Many downstream paths are still 750 MHz or 862 MHz. With an increase to full DOCSIS 3.1 1218 MHz spectrum and with video reclamation, this could provide a 200% to 400% increase in data capacity. With DOCSIS 4.0, there can be an additional 40% to 60% increase. With these numbers multiplied, the downstream has a 10x growth potential in data capacity.

The upstream spectrum is defined by a 42 MHz return path with about 100 Mbps. An upgrade to 204 MHz would be a 1400% increase and DOCSIS 4.0 could more than double that again. So, the return path has considerably more than 10x growth to go.



**Figure 11 – DOCSIS Plans for Spectrum Increases**

Then there is segmentation. A 500 HHP N+5 plant could be segmented down to 50 HHP N+0. These are approximate numbers as each plant is unique. However, this represents another 10x increase in capacity.

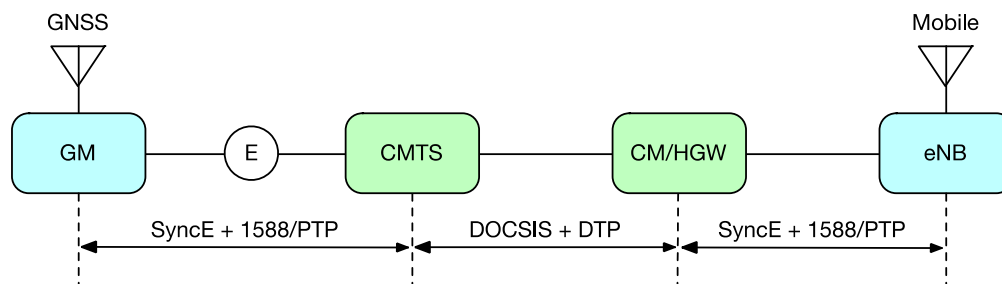
Combining the 10x growth available through spectrum and the 10x growth available through segmentation, [there is 100x growth potential available in the DOCSIS HFC network](#).

CBRS small cells are on the order of 100 to 200 Mbps, depending upon spectrum usage and distance. Cable Modems today are 100 Mbps to 1 Gbps in the downstream. So, roughly, another small cell load is like another CM load, or less, on the network.

DOCSIS could even be extended to 3 GHz one day in the downstream using distributed gain amplifiers (DGA). This was first discussed in [15]. As an alternative to extending the downstream frequency range, DOCSIS 4.0 permits the use of full-duplex technology which would allow up to a 5 Gbps in the 5 to 684 MHz return spectrum to co-exist with the 10 Gbps in the 108 to 1218 MHz forward spectrum downstream spectrum. This technology was laid out in the following papers [16][17][18][19][20][21][22].

So, yes, DOCSIS does have enough bandwidth to handle small cell distribution.

### 4.3. SYNC – 1588 and SyncE Support



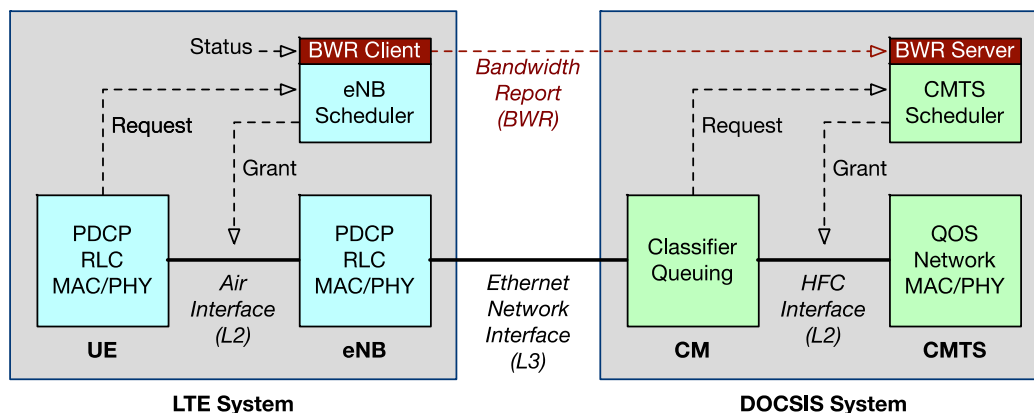
**Figure 12 – SYNC Operating Model**

In addition to the three basics of location, power, and backhaul, small cells that are located indoors and do not have access to the GNSS network, will need timing support. This is often done by having a synchronized wireline network using Synchronous Ethernet and IEEE-1588. Support for these two protocols are being designed into the DOCSIS protocol and product. The approach is shown in Figure 12.

DOCSIS is a highly accurate synchronous network. DOCSIS 3.1 defines a method to derive the DOCSIS timestamp from a Precision Time Protocol (PTP) slave port. DTP (DOCSIS Time Protocol) is used to measure the time difference across the DOCSIS network. The CM implements a PTP master clock and uses this time difference and the DOCSIS timestamp to regenerate an accurate PTP clock. This has been standardized at CableLabs [15] [24] and is described further in [25][26][27].

LTE small cells at 700 MHz are frequency division duplex (FDD) based and can use ranging to work around timing. This allows LTE FDD small cells to not have to worry about network timing. CBRS small cells at 3.5 MHz typically uses TDD (Time Division Duplex) which requires tight timing synchronization of +/- 1500 ns from radio to radio in order to prevent the radios from stepping on each other's transmissions.


### 4.4. LLX – Low Latency Xhaul



**Figure 13 – LLX Operating Model**

As mobile networks move to 5G, there is an interest in low latency operation. DOCSIS has a new protocol called LLX that pipelines the requests from the eNB/gNB to the CMTS scheduler so that packets do not have to re-request for bandwidth when entering the DOCSIS system. This will allow the DOCSIS system to have approximately 2 ms of equivalent latency. This is illustrated in Figure 13. This has been standardized at CableLabs [28] and is described further in [29][30][31][32][33].

## 5. Conclusion



**Figure 14 – Macrocells and Small Cells in Deployment**

We started this paper with a deployment model from the City of Danville and we are ending with another diagram in Figure 14 also from the city of Danville, CA [1] . This diagram clearly shows that small cells will be located on poles, in buildings and in homes. Today’s foreground macrocell based LTE network will become tomorrow’s background network to fill in the gaps in the new 5G small cell architecture.

Here is a summary of the important formulas:

$$\#SC \text{ per MC} = \pi/2 (R/r)^2 \quad \text{per formula (3)}$$

$$\#SC \text{ per FN} = \pi/4 (S/r)^2 (M + 1)^2 \quad \text{per formula (11)}$$

where:

$R$  = larger radius of macrocell or node  
 $r$  = smaller radius of small cell  
 $S$  = average coax span between actives  
 $M$  = number of amplifiers in an N+M cascade

Here is a summary of the important points raised in this white paper:

- Today’s cable operators are tomorrow’s mobile operators
- Behind every great wireless network is a great wireline network
- The number of radios for CBRS can be 100x that of an LTE macrocell, and for mmWave, the number of radios could be 100,000x that of an LTE macrocell.
- This dramatically impacts deployment economics and makes the existing HFC plant and interesting choice for mobile backhaul.
- The number of small cells in an HFC plant is proportional to the square of the span count and the square of the ratio of span length to cell radius.
- DOCSIS has a 100x growth potential in downstream and upstream bandwidth
- The DOCSIS HFC network is a viable backhaul/midhaul network that can meet the location, bandwidth, powering, timing, and latency requirements of a small cell network.

# Abbreviations

|        |                                                |
|--------|------------------------------------------------|
| 5G     | Fifth Generation Mobile Network                |
| 5GC    | 5G Core                                        |
| BAU    | Business as Usual                              |
| CBRS   | Citizens Broadband Radio Service               |
| CM     | Cable Modem (DOCSIS)                           |
| CMTS   | Cable Modem Termination System (DOCSIS)        |
| CU     | Centralized Unit                               |
| DAA    | Distributed Access Architecture                |
| DGA    | Distributed Gain Amplifier                     |
| DOCSIS | Data over Cable System Interface Specification |
| DTP    | DOCSIS Time Protocol                           |
| DU     | Distributed Unit                               |
| EIRP   | Effective Isotropic Radiated Power             |
| eNB    | Evolved Node B (LTE)                           |
| FDD    | Frequency Division Duplex                      |
| GAA    | General Authorized Access                      |
| gNB    | Next generation NodeB                          |
| GPRS   | General Packet Radio Service                   |
| GTP    | GPRS Tunneling Protocol                        |
| HFC    | hybrid fiber-coax                              |
| HHP    | households passed                              |
| IP     | Internet Protocol                              |
| LLX    | Low Latency Xhaul                              |
| MVNO   | Mobile Virtual Network Operator                |
| LTE    | Long Term Evolution                            |
| N+M    | Node plus M amplifiers in an HFC network       |
| PAL    | Priority Access License                        |
| PDCP   | Packet Data Convergence Protocol               |
| PTP    | Precision Time Protocol                        |
| RF     | Radio Frequency                                |
| RLC    | Radio Link Control                             |
| RRC    | Radio Resource Control                         |
| RU     | Radio Unit                                     |
| SAS    | Spectrum Access System                         |
| SHO    | Super High Output                              |
| TDD    | Time Division Duplex                           |
| XHAUL  | Backhaul, midhaul, or fronthaul                |
| UE     | User Equipment                                 |

# Bibliography & References

- [1] John T. Chapman, “DOCSIS Remote PHY”, *SCTE Cable-Tec Expo*, Oct 2013.
- [2] John T. Chapman, “Remote PHY for Converged DOCSIS, Video and OOB”, NCTA Technical Conference, Jun, 2014. [[link](#)]
- [3] “Small Cell Wireless Facilities Fact Sheet”, City of Danville, Nov, 2018. [[link](#)]
- [4] Commscope Ruckess Q910 Outdoor Data Sheet. [[link](#)]
- [5] “U.S. Telecom and Cable & Satellite Marketing Deck”, MoffettNathanson Research, Apr, 2020, page 120.
- [6] Kyung Mun, “CBRS White Paper: CBRS: New Shared Spectrum Enables Flexible Indoor and Outdoor Mobile Solutions and New Business Models”, Mobile Experts LLC, 2017 [[link](#)]
- [7] Jennifer Andreoli-Fang, John T Chapman, et. al., “Cable and Mobile Convergence – A Vision From the Cable Communities Around the World”, *SCTE-Tec Expo Fall Technical Forum*, Oct, 2020. [[link](#)]
- [8] Broadband Technology Report, “Lindsay Broadband, Accelleran team with Cogeco for HFC plant-powered small cell field trial”, Jul, 2020. [[link](#)]
- [9] Damian Poltz, “HFC and Wireless – Cable’s Convergence Advantage”, CableLabs Summer Conference, Aug, 2019
- [10] Heavy Reading, “White Paper – Cable’s Value Proposition for Small Cells”, Dec, 2015 [[link](#)]
- [11] Fierce Wireless, “Sprint inks small cell deal with Cox but remains silent on MVNO front”, Jan, 2018. [[link](#)]
- [12] John T Chapman, “Blog - Mobile Xhaul Over DOCSIS Delivers Faster Time to Market at a Lower Cost Than Building a New Fiber Plant”, Cisco Service Provider Blog Site, Sep, 2019. [[link](#)]
- [13] Jennifer Andreoli-Fang, John T Chapman, Tong Liu, Damian Poltz, “Blueprint for Mobile Xhaul over DOCSIS,” *SCTE Cable-Tec Expo*, Sep, 2019. [[link](#)]
- [14] “DOCSIS Physical Layer Specification”, CM-SP-PHY, CableLabs [[link](#)]
- [15] John T. Chapman, Hang Jin, Thushara Hewavithana; Rainer Hillermeier, “Blueprint for 3 GHz, 25 Gbps DOCSIS,” *SCTE Cable-Tec Expo Fall Technical Forum*, Sep, 2019 [[link](#)]
- [16] John T. Chapman, Hang Jin, “Full Duplex DOCSIS”, *SCTE/NCTA Spring Technical Forum*, May, 2016. [[link](#)]
- [17] Tong Liu, John T. Chapman, Hang Jin, “Interference-Aware Spectrum Resource Scheduling for FDX DOCSIS”, *SCTE Journal of Network Operations*, Vol 1, No 2, Sept, 2016. [[link](#)]

- [18] Hang Jin & John T Chapman, “Echo Cancellation Techniques for Supporting Full Duplex DOCSIS.”, *SCTE Cable-Tec Expo Fall Technical Forum*, October, 2017. [[link](#)]
- [19] John T Chapman, Hang Jin, “FDX DOCSIS Line Extender: Deploying FDX DOCSIS Beyond N+0”, *SCTE Cable-Tec Expo Fall Technical Forum*, Oct, 2018 [[link](#)]
- [20] Hang Jin, John T. Chapman, “FDX Amplifier for Supporting N+M Network”, *SCTE Cable-Tec Expo Fall Technical Forum*, Sep, 2019. [[link](#)]
- [21] Tong Liu, “Characterization of Spectrum Resource Scheduling in FDX DOCSIS”, *SCTE Cable-Tec Expo Fall Technical Forum*, Oct, 2018. [[link](#)]
- [22] Tong Liu, “Interference Group Discovery for FDX DOCSIS”, *SCTE Cable-Tec Expo Fall Technical Forum*, Oct 2017. [[link](#)]
- [23] “Synchronization Techniques for DOCSIS Technology Specification,” CM-SP-SYNC, CableLabs. [[link](#)]
- [24] “DOCSIS MAC and Upper Layer Protocols Interface Specification”, CM-SP-MULPI, CableLabs. [[link](#)]
- [25] Elias Chavarria Reyes, John T. Chapman, “How the DOCSIS Time Protocol makes the SYNC Specification Tick,” SCTE Cable-Tec Expo, Oct, 2020. [[link](#)]
- [26] Jennifer Andreoli-Fang, John T. Chapman, “Mobile Backhaul Synchronization Architecture,” SCTE Fall Technical Forum, October, 2017. [[link](#)]
- [27] John T. Chapman, et. al., “The DOCSIS Timing Protocol (DTP), Generating precision timing services from a DOCSIS system,” INTX/SCTE Spring Technical Forum, 2011. [[link](#)]
- [28] “Low Latency Mobile Xhaul over DOCSIS Technology,” CM-SP-LLX, CableLabs. [[link](#)]
- [29] John T. Chapman, Jennifer Andreoli-Fang, Michel Chavin, Elias Chavarria Reyes, Zheng Lu, Dantong Liu, Joey Padden, Alon Bernstein, “Low latency techniques for mobile backhaul over DOCSIS,” Proc. of IEEE Wireless Communication and Networking Conference (WCNC), Barcelona, April 2018. [[link](#)]
- [30] John T. Chapman, Jennifer Andreoli-Fang, Michel Chavin, Elias Chavarria Reyes, Zheng Lu, Dantong Liu, Joey Padden, Alon Bernstein, “Low latency techniques for mobile backhaul over DOCSIS,” Proc. of IEEE Wireless Communication and Networking Conference (WCNC), Barcelona, April 2018. [[link](#)]
- [31] John T. Chapman, Jennifer Andreoli-Fang, “Low Latency Techniques for Mobile Backhaul over DOCSIS,” SCTE Fall Technical Forum, October, 2017. [[link](#)]
- [32] Jennifer Andreoli-Fang, John T. Chapman, “Mobile-aware scheduling for low latency backhaul over DOCSIS,” Proc. of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), Montreal, Oct 2017. [[link](#)]
- [33] Jennifer Andreoli-Fang, John T. Chapman, “Latency reduction for mobile backhaul over DOCSIS through pipelining,” Proc. of IEEE Globecom, Singapore, Dec 2017. [[link](#)]



# Content Aware Video Streaming

A Technical Paper prepared for SCTE•ISBE by

**Srinath V Ramaswamy**

Principal Solution Architect

Comcast

1717 Arch Street, Philadelphia, PA 19103

(215) 286 5444

srinath\_ramaswamy@cable.comcast.com

# Table of Contents

| <b>Title</b>                                                  | <b>Page Number</b> |
|---------------------------------------------------------------|--------------------|
| 1. Introduction .....                                         | 3                  |
| 2. IP Video Architecture .....                                | 3                  |
| 3. Segment Video Analysis .....                               | 4                  |
| 3.1. Enhanced IP Video Architecture with Video Analysis ..... | 7                  |
| 4. Carriage of Optional Download Information .....            | 7                  |
| 5. Benefits .....                                             | 8                  |
| 6. IP Video Network Bandwidth Utilization .....               | 8                  |
| 7. Conclusion .....                                           | 9                  |
| Abbreviations.....                                            | 9                  |
| Bibliography & References .....                               | 10                 |

## List of Figures

| <b>Title</b>                                            | <b>Page Number</b> |
|---------------------------------------------------------|--------------------|
| Figure 1 – IP Video Architecture .....                  | 4                  |
| Figure 2 – IP Video Enhanced Architecture.....          | 7                  |
| Figure 3 – IP Video Network Bandwidth Utilization ..... | 9                  |

## List of Tables

| <b>Title</b>                                                     | <b>Page Number</b> |
|------------------------------------------------------------------|--------------------|
| Table 1 – Constant AVC Video Bitrate Ladder .....                | 4                  |
| Table 2 – P/B Frame High Activity and Detail.....                | 6                  |
| Table 3 – Bitrate Ladder after Video Analysis for Segment n..... | 6                  |

## 1. Introduction


Streaming video represents nearly 80% of all traffic carried over the internet. It's growing by 22% each year [1]. Technological evolution in video compression, from MPEG-4/AVC [2] to HEVC [3], paired with techniques like Adaptive Bit Rate (ABR), to “right size” a video segment into multiple bit rates, represent some of the core building blocks that contributed to the huge marketplace success that is streamed video. However, a third building block is becoming necessary, to go a step further in both an improved, adaptive video experience, at the device end, and optimized bandwidth usage, to lessen network strain. Specifically, and considering the remarkable and mushrooming growth of IP video streaming, there exists a growing need for systems to leverage Content Aware Video Streaming solutions, like the one discussed in this paper.

In IP video streaming, video is typically generated at several bitrates, in order to accommodate the varying network bandwidth available between IP Video consumer devices and content delivery networks (CDNs) and varying consumer devices. Consumer devices will typically try to download the highest bitrate video possible that the device can support, with the intent of providing higher video quality to the viewer. Higher bitrate video segments for certain content do not necessarily translate to higher video quality -- if there is any improvement at all, it might be marginal. Acquiring higher bitrate segments when it does not really enhance user video viewing experience results in wasted network bandwidth. This wastage can be avoided by having the devices only download segments that would make a difference in video quality for the viewer. For the devices to do a quality-based download decision, they would need to be provided with info on the content in the segment. In this paper, an algorithm that determines if a segment is going to enhance viewer video experience is presented. This algorithm is based on video analysis without a full video decode of frames in the segment. This algorithm would need to be implemented at the content origination end and the generated info transmitted to consumer devices.

With the content information communicated in each segment, devices can then make a network bandwidth and quality based decision on the next segment to be downloaded. This would minimize network bandwidth usage and at the same time providing the highest video quality. The algorithm presented supports processing of live linear stream segments and would minimally impact the content generation workflow and therefore scalable. Included in this paper are results showing network bandwidth savings realized by this algorithm. This algorithm could also be used to reduce storage needs by only storing segments from a bitrate ladder that enhances video quality on various devices.

## 2. IP Video Architecture

At a high level, the IP Video architecture is as shown in the Figure 1 below. In this figure, the compressed or uncompressed video stream is transcoded into several streams (variants), each at different fixed bitrate, to accommodate the varying network bandwidth available between IP Video consumer devices and the content delivery network. The transcoded streams might be encoded using one of the MPEG compression codecs (MPEG-2[4], MPEG-4/AVC, HEVC). This is then packaged into one of the many streaming formats, for example MPEG-DASH [5], or HLS [6]. These are then placed on the Origin Server and IP CDN for delivery to IP Video clients. This would apply to both linear and on demand IP video.



**Figure 1 – IP Video Architecture**

### 3. Segment Video Analysis

The IP Video variant streams shown in the Figure 1 are typically split into segments of a few seconds duration, and these segments across the variants are assumed to be time-aligned. The IP Video variant streams are part of a constant bitrate ladder as shown in Table 1, below:

**Table 1 – Constant AVC Video Bitrate Ladder**

| Resolution | Video Bitrate (Kbps) |
|------------|----------------------|
| 1920x1080  | 6000                 |
| 1920x1080  | 4000                 |
| 1280x720   | 3000                 |
| 1280x720   | 2000                 |
| 720x480    | 1600                 |
| 640x480    | 1000                 |
| 512x384    | 700                  |
| 384x288    | 400                  |
| 384x288    | 300                  |

An IP Video consumer device traverses up and down this ladder, based on varying network conditions between it and the CDN. During stable network conditions the device might be playing segments from the 6 Mbps video stream but for some of these segments it might be of the same quality as the 4 Mbps segments, therefore the device could downshift to the 4 Mbps segments for those segments and then go back to the 6 Mbps segments when there is an actual gain in the video quality by doing so. By taking this approach network bandwidth usage is reduced. Similarly, in another scenario when the network condition improves after a degradation and the device sees that the network can sustain a higher bitrate, the device which is currently playing a 4 Mbps video stream would try to display a high quality video by moving up the bitrate ladder. In this case, it would try to download and play a 6 Mbps higher bitrate variant segment, even though the gain in the video quality by doing so may be minimal. So in both these scenarios, during the download process, if the device is told that downloading the 6 Mbps segment would result in a minimal video quality gain, then it could avoid the unnecessary download, which saves network bandwidth. Therefore, there is a need for an approach that can detect if a segment would enhance video quality and communicate that to IP Video devices. Also, with this information, the device could look at downloading the next higher bitrate variant segment that really would improve the user experience.

Similarly, when the network condition is worsening and the device is moving down the bitrate ladder, and if it is told that the next lower bitrate stream does not enhance quality, then it could select the one below it in the bitrate ladder, again saving network bandwidth.

### Video Analysis Algorithm Details

In this section, an encoded video analysis algorithm is detailed that determines if an IP video segment does not enhance video quality and is therefore optional to download. The results of this algorithm are then communicated to the IP video device, which would utilize it to determine the segments to download, in order to save network bandwidth and enhance video quality.

The algorithm analyzes all encoded video frames in a segment to determine if it is optional to download the segment. This is done in the following manner:

1. The compressed bitstream in the segment is first entropy decoded and inverse quantized to access the many frames in the segment and its compression details.
2. The frame type (Intra [I], Predictive [P], Bidirectional [B]) is determined for each frame in the segment. The frame type determines how that frame is going to be analyzed to determine its contribution to the segment relevancy.
3. If the frame is an I type, we determine the frame's compression ratio (defined as the number of bytes to represent the frame / the number of bytes to represent the compressed frame). If the compression ratio is less than a predetermined threshold, then this frame is marked as a frame with high detail and activity. If the compression ratio is low, it would mean that there are a lot of details in this frame and less redundancy.
4. If the frame is a P or B type, we again determine the frame's compression ratio. As with I-frames, if the compression ratio is less than a predetermined threshold then this frame is marked as one with high detail and activity. If the compression ratio is higher, then some of its compression attributes are analyzed to determine if it would still qualify as a frame with high detail and activity. The compression attributes that are analyzed are listed below:
  - Number of skipped Macroblocks/Coding Unit (CU): The number of skipped macroblocks in the frame is determined. If it exceeds a certain threshold, then it is counted towards the determination of a high level of detail and activity and is marked as a High Skipped MB/CU frame. A frame with a lot of skipped macroblocks indicates that there isn't much motion between this frame and its previous frames, used for motion estimation. A macroblock is the basic processing unit in AVC/MPEG-2, and the CU, in HEVC, is where the prediction type is decided, as well as which frames are skipped when their motion vector is zero and coefficients are zero.
  - Partition (Prediction) block size: A prediction block is the block split from a macroblock or coding unit that is used for motion estimation. This attribute determines the size of the prediction block that is the most widely used in this frame. If it exceeds a certain threshold, then it is marked as a Large Partition Block frame. A comparatively small size indicates a lot of details in the frame.
  - Inter coded blocks count: Determines the number of inter-coded blocks in frame. If the number of inter-coded blocks count exceeds a threshold, then it is indicated as High Inter-Coded Block Count frame. A low number of inter-coded blocks implies that there is a scene change, therefore unable to predict from reference frames.
  - Motion vector component measure: Determines motion vectors for each inter-prediction partition block in the frame and determines the 90% trimmed standard deviation of the motion vectors' horizontal and vertical components. The 90% trimmed standard deviation of the components is to ignore extreme values to give a better measure and is determined by ignoring the lower 5% and top 5% of the values. These vectors are

encoded differentially, with respect to predicted values from nearby vectors. If greater than a certain threshold, it would be considered a high MV measure frame. Large variations or standard deviations would indicate a lot of activity.

Note that the thresholds used for these compression attributes for P frames are different from those used for B frames, and these thresholds for P and B frames also factor in video frame rates and resolution.

The information gathered by analyzing these compression attributes is then used to determine if this P or B frame is a high detail and activity frame, as shown in Table 2, below:

**Table 2 – P/B Frame High Activity and Detail**

| <b>High Skipped MB/CU</b> | <b>Large Partition Blocks</b> | <b>High Inter Coded Blocks Count</b> | <b>High MV Measure</b> | <b>High Activity/ Detail?</b> |
|---------------------------|-------------------------------|--------------------------------------|------------------------|-------------------------------|
| False                     | False                         | False                                | True                   | True                          |
| True                      | N/A                           | N/A                                  | N/A                    | False                         |
| False                     | True                          | False                                | True                   | True                          |
| False                     | True                          | True                                 | True                   | False                         |
| False                     | True                          | True                                 | False                  | False                         |
| False                     | False                         | True                                 | True                   | True                          |
| False                     | False                         | True                                 | False                  | False                         |
| False                     | False                         | False                                | True                   | False                         |

Once all the frames in a segment have been analyzed using the approach detailed above, the percentage of high activity/ detail frames in the segment is determined. If this percentage is below a certain threshold value, then it is marked as a segment for optional download. So, when an IP video consumer device is trying to display a higher quality video by moving up the bitrate ladder and encounters this optional download segment, it would not download it, because there is no appreciable gain in video quality by doing so. It would, however, consider downloading a higher bitrate segment that is not marked as an optional download.

Table 3 shows the resultant bitrate ladder after the video for segment “n” has been analyzed, using the algorithm detailed in this paper:

**Table 3 – Bitrate Ladder after Video Analysis for Segment n**

| <b>Resolution</b> | <b>Video Bitrate (Kbps)</b> | <b>Download for Higher Quality?</b> |
|-------------------|-----------------------------|-------------------------------------|
| 1920x1080         | 6000                        | No                                  |
| 1920x1080         | 4000                        | Yes                                 |
| 1280x720          | 3000                        | Yes                                 |
| 1280x720          | 2000                        | Yes                                 |
| 720x480           | 1600                        | No                                  |
| 640x480           | 1000                        | Yes                                 |
| 512x384           | 700                         | Yes                                 |

| Resolution | Video Bitrate (Kbps) | Download for Higher Quality? |
|------------|----------------------|------------------------------|
| 384x288    | 400                  | No                           |
| 384x288    | 300                  | Yes                          |

This video segment analysis does not need to be done for all the video segments across the variants -- only those necessary to determine a high level of activity and detail for each segment time duration. For example, each segment duration could be 4 seconds across the variants. Since the encoded video analysis is done without a full decode, this algorithm does not add latency, in the case of linear video, or increase encoding time, in the case of non-linear video.

### 3.1. Enhanced IP Video Architecture with Video Analysis

An enhanced architecture that utilizes the algorithm detailed in this paper is shown in the Figure 2 below. The key difference between this and the architecture shown in Figure 1 is the addition of a functional block that processes the IP video segments and generates content-specific information sent to the client.




Figure 2 – IP Video Enhanced Architecture

## 4. Carriage of Optional Download Information

There are multiple approaches to send the optional segment download information to IP Video consumer devices.

One approach would be to signal those segments that are optional to download in the manifest files. For example, in the case of DASH, this would happen in the Representation element; in HLS this would happen in the Media Playlist. In the case of live linear, as the segments are generated, the optional

download information is determined and included in the updated manifest files. When signaled in the manifest, the consumer device could base its next download decision on this value.

Another approach would be to remove the segment that is optional to download from a Representation, in a DASH manifest file, or, in the case of HLS, to remove the media file URL from the Media Playlist.

Potentially one other approach would be to indicate the optional segment download information via fields in the segment. For example, in user data fields in the MPEG-2 Transport packet headers, or ISOBMFF m4s file box structures.

Please note that none of these are currently standardized.

## **5. Benefits**

There are many benefits reaped by utilizing the Content Aware streaming solution described in this paper. Primarily by not downloading higher bitrate segments that does not necessarily provide a gain in video quality streaming network bandwidth usage is reduced. Lab studies have shown a reduction of 20%.

Since this algorithm identifies segments that do not contribute to higher video quality, one approach would be to not store those segments and instead point the consumer device to an alternate segment that has the same quality as this one. This would save storage in applications like Cloud DVR, Video On Demand and CDN. Another benefit from this is that since the segments are of smaller size during certain times, the downloads will be faster.


The algorithm detailed in this paper does not do a full decode of the compressed bits in a segment but does partial decode resulting a fast video quality assessment.

The intent of this solution is to not impact the existing transcoders and packagers but execute this algorithm on the content generated by packagers before they are made available to IP video consumer devices.

## **6. IP Video Network Bandwidth Utilization**

Figure 3 shows instantaneous video bitrate when consumer device is downloading segments with and without using the solution detailed in this paper. The black curve plots the video bitrate when segments are downloaded without looking at the segment quality whereas the blue curve shows the bitrate when the consumer device downloads segments based on its quality. So, for example at Time 7 the device downloads a segment from a lower bitrate stream that has the same video quality as the segment downloaded at the higher bitrate indicated on the black curve. It is observed that the video quality is the same with the two type of downloads.





**Figure 3 – IP Video Network Bandwidth Utilization**

## 7. Conclusion

IP Video Streaming is widespread and consuming valuable network bandwidth – as in, almost 80% of all internet traffic. This traffic is expected to grow by 22% yearly. Considering this growth, it is important for systems to leverage Content Aware Video Streaming solutions such as this to optimize network bandwidth usage. In this paper an algorithm is presented that determines if an IP video segment is going to enhance viewer video experience. This algorithm, which analyzes encoded video without a full decode, in linear IP video. Also included are approaches to communicate an IP video segment’s relevance to enhance viewer video quality to IP video devices. In addition to optimizing network bandwidth usage and adapting to changing network conditions, it also reduces storage requirements and enhances consumer video viewing quality.

## Abbreviations

|          |                                                              |
|----------|--------------------------------------------------------------|
| ABR      | Adaptive Bitrate                                             |
| AVC      | Advanced Video Coding                                        |
| CDN      | Content Delivery Network                                     |
| DASH     | Dynamic Adaptive Streaming over HTTP                         |
| HEVC     | High Efficiency Video Coding                                 |
| HLS      | HTTP Live Streaming                                          |
| IP       | Internet Protocol                                            |
| ISO-BMFF | International Standard Organization – Base Media File Format |

|      |                              |
|------|------------------------------|
| MPEG | Motion Picture Experts Group |
| URL  | Uniform Resource Locator     |

## Bibliography & References

[1] [https://www.cisco.com/c/dam/m/en\\_us/solutions/service-provider/vni-forecast-highlights/pdf/United\\_States\\_2022\\_Forecast\\_Highlights.pdf](https://www.cisco.com/c/dam/m/en_us/solutions/service-provider/vni-forecast-highlights/pdf/United_States_2022_Forecast_Highlights.pdf)

[2] H.264, Coding of Audio-Visual Objects — Part 10: Advanced Video Coding (AVC), ISO/IEC 14496-10:2010.

[3] H.265, High efficiency coding and media delivery in heterogeneous environments — Part 2: High efficiency video coding, ISO/IEC 23008-2:2013

[4] H.262, Generic Coding of Moving Pictures and Associated Audio Information (H.262), ISO/IEC 13818-2:2013.

[5] Dynamic Adaptive Streaming over HTTP (DASH), ISO/IEC 23009-1:2014

[6] R. P. Pantos, editor, HTTP Live Streaming, <https://tools.ietf.org/html/rfc8216>

# **Convolutional Neural Networks for Proactive Network Management**

## **Developing Machine Learning Models to Detect and Classify Impairments in D3.1 OFDM Channels**

A Technical Paper prepared for SCTE•ISBE by

**Jude Ferreira**

Principal Data Scientist

Comcast

215.286.4070

jude\_ferreira@cable.comcast.com

**Maher Harb**

Director, Data Science

Comcast

267.260.1846

maher\_harb@comcast.com

**Karthik Subramanya**

Research Engineer

Comcast

267.260.2289

karthik\_subramanya@comcast.com

**Bryan Santangelo**

Executive Director, Data Eng and Science

Comcast

918.640.8936

bryan\_santangelo@cable.comcast.com

**Dan Rice**

Vice President, HFC Architecture

Comcast

720.512.3730

daniel\_rice4@comcast.com

# Table of Contents

| Title                                                                | Page Number |
|----------------------------------------------------------------------|-------------|
| 1. Introduction .....                                                | 3           |
| 2. Data Collection: OFDM Receive Modulation Error Ratio (RxMER)..... | 4           |
| 3. Labels and Supervised Learning .....                              | 4           |
| 4. Data Preprocessing .....                                          | 5           |
| 5. Model Evaluation .....                                            | 6           |
| 6. Convolutional Neural Networks(CNNs).....                          | 8           |
| 7. Modeling .....                                                    | 9           |
| 8. Results .....                                                     | 10          |
| 9. Machine Learning Pipeline .....                                   | 11          |
| 9.1. Data Lake.....                                                  | 12          |
| 9.2. Compute Engine.....                                             | 12          |
| 9.3. Integration layer:.....                                         | 13          |
| 10. Conclusion/Next Steps .....                                      | 13          |
| Abbreviations.....                                                   | 14          |
| Bibliography & References .....                                      | 15          |

## List of Figures

| Title                                                                                             | Page Number |
|---------------------------------------------------------------------------------------------------|-------------|
| Figure 1. Different types of impairments seen in D3.1 OFDM Channels .....                         | 3           |
| Figure 2. Screenshot of the RxMER Pattern Labeling UI .....                                       | 5           |
| Figure 3. Number of samples by label (impairment type).....                                       | 6           |
| Figure 4. Distribution of number of impairments/samples .....                                     | 7           |
| Figure 5. Convolution Neural Network (CNN) components.....                                        | 8           |
| Figure 6. Network Architecture CNN model that had best performance on validation dataset [2]..... | 10          |
| Figure 7. ROC Curves .....                                                                        | 11          |
| Figure 8. Machine Learning Pipeline Layers.....                                                   | 12          |

## List of Tables


| Title                                                                                                         | Page Number |
|---------------------------------------------------------------------------------------------------------------|-------------|
| Table 1. Hyperparameters, Ranges evaluated during training .....                                              | 9           |
| Table 2. Subset Accuracy, Hamming Loss for models.....                                                        | 10          |
| Table 3. CNN model with average MER padding - Individual classes performance metrics on validation data ..... | 11          |

## 1. Introduction

The signal quality on HFC networks can degrade over time, from an impairment perspective, if not proactively maintained. Comcast manages hundreds of thousands of miles of network throughout the world in which we experience a wide array of conditions that can degrade the performance of the network. From connections loosening and cracks forming, to lines getting cut, destructive energy and signal impediments are part of maintaining modern networks. Early detection, mitigation, and routing fix agents efficiently, to the right location, can not only improve customer experience but also enable operational efficiency.

Adaptive Profile Management Application (PMA) systems continue to be deployed to manage downstream and upstream network capacity and network stability. However, and perhaps ironically, PMA systems can mask the degradation of the network, as an inherent function of the optimization and mitigation process. It therefore has become increasingly important to develop systems that can support automated Proactive Network Maintenance (PNM) to reduce the impact of network impairments on customer experience and enable the highest possible capacity and performance.

In this regard, Comcast has invested heavily in data platforms and data science functions across organizations, to become more data driven and to incorporate Machine Learning (ML) approaches into the network. In this paper, we will describe the use of Convolutional Neural Networks (CNNs) to identify network impairments within DOCSIS 3.1 (D3.1) channels with a high degree of accuracy.



**Figure 1. Different types of impairments seen in D3.1 OFDM Channels**

It is beneficial to classify the various network impairments (shown in Figure 1) as they may warrant different responses from techs in the field to enable fastest possible Mean Time To Repair (MTTR). In

addition, clustering of these impairments across geographic locations and network topology may be exploited to identify the root cause impacting multiple customers that share common points in the network. Note that the latter requires a second layer model to be built on top of the classification model described in the paper.

The model we describe improves on the rule-based approaches currently being used to identify Mobile Wireless Ingress and Sweep Generator patterns. Notifications from the current rule-based model for detecting Mobile Wireless Ingress are sent across a notification bus to other Comcast OSS tools, to ensure that technicians are dispatched to the right hubs, network segments, and homes to remediate issues. The rule-based approach also provides a baseline for evaluation for the ML approaches. We also developed a real time version of the algorithm that techs can use to check that issues have been fixed after remediation. The same workflow described in this paper could also be used with alternate CNN-based models.

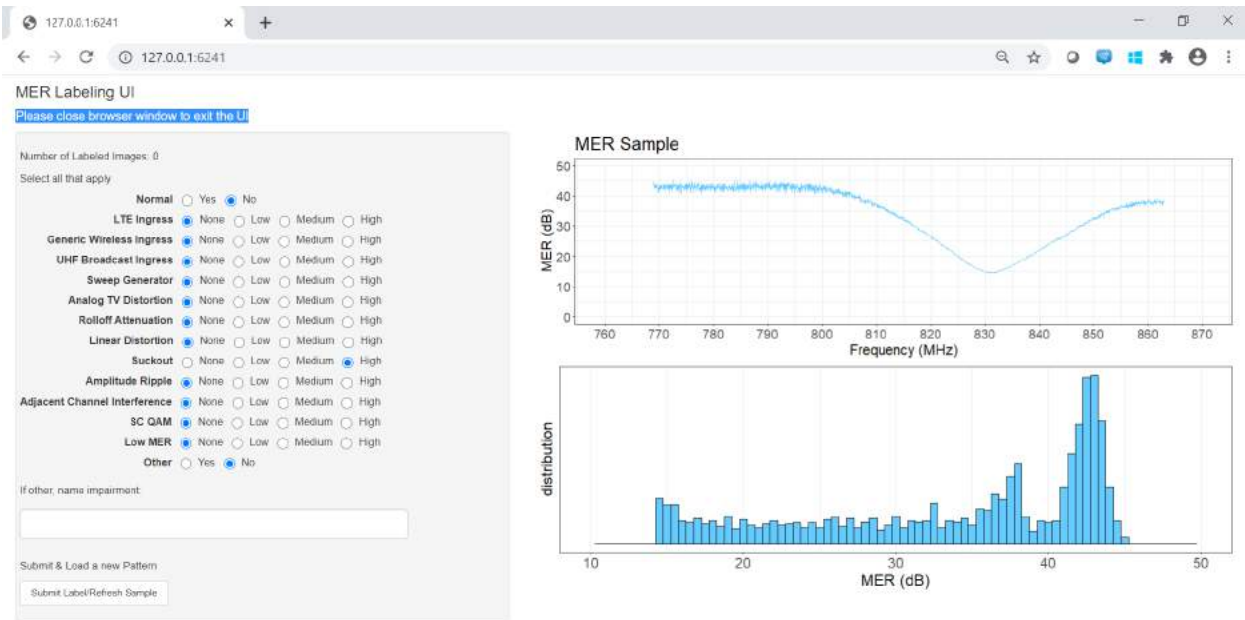
## **2. Data Collection: OFDM Receive Modulation Error Ratio (RxMER)**

The underlying data for the model comes from a data collection platform, which acquires telemetry from cable modems and CMTSs for various performance-related measures, as well as other characteristics, such as make, model, hardware and software versions. The API platform runs both on pre-determined intervals as well as on-demand to collect data from the entire access network. Real time use cases are typically focused to specific groups of cable modems or interfaces with higher periodic rates. The collection platform also provides synchronous and asynchronous API requests so that response from cable modems and CMTSs can be returned back to the consumer, sent on a message bus for multiple consumers, and captured into our data lake.

Comcast supports millions of D3.1 devices. The data collection captures high resolution Receive Modulation Error Rate (RxMER) data from these devices at regular intervals. The methods described to identify various impairments for OFDM Channels in this paper are based on this high resolution RxMER measure per subcarrier (as shown in Figure 1). Modulation Error Ratio is an important measure for consideration, as it not only picks up on core signal-to-noise (SNR) characteristics but also all signal imperfections. Therefore, high resolution (high frequency and per-OFDM subcarrier) captures of RxMER is the primary measure and focus for characterization of impairments.

## **3. Labels and Supervised Learning**

Methods for training models with well labeled data sets are well established and vastly varietal. While initial attempts focused on non-supervised machine learning methods, such as clustering, to separate out impairments into similar groups, the results were not promising. Thus, we recognized the need to label a set of training data to be applied to supervised machine learning methods. Labeling can be labor-intensive activity that, in this paradigm, would also rely on subject matter experts to examine and classify the impairments based on their expertise. In order to make the labeling process robust, we developed a Pattern Labeling user interface (UI) that makes RxMER samples available to labelers. To help capture impairments, and given that a majority of RxMER samples do not have impairments, the sampling strategy focused on capturing samples with high variance over OFDM subcarriers. When presented, the labelers can examine and submit their assessment for the impairments within a few mouse clicks.



**Figure 2. Screenshot of the RxMER Pattern Labeling UI**

Labels are a critical component to building a good performing classification model. Given the significant value of enhancing capacity and improving customer experience by being able to identify and characterize specific impairments, we aimed for a crowdsourcing approach involving field technicians and other SMEs to generate labeled data using this UI. Once initial models are developed, we plan to pre-populate the Labeling UI with the impairments predicted by the different models to make the process more efficient.

D3.1 OFDM Channels are configured to extend from 24 to 192 MHz and often placed in the highest spectrum of the Hybrid Fiber Copper (HFC) network above the video and D3.0 channels. In many service groups, some portion of the OFDM channel is outside the HFC design, in what is often referred to as the “roll-off” spectrum. Given the size of the channels and where the OFDM channel may be located, it is possible that some cable modems could experience multiple impairments. Thus, the Labeling UI allows a sample to be labeled with multiple impairments. Also, when impairments exist, the level of severity (low, medium, high) needs to be specified. Future versions may allow for more nuanced severities. Note that characterizing severity will likely increase the number of training samples needed to build good models. The labeling effort is designed such that each sample would get multiple responses from different experts, to resolve conflicting labels and to build confidence in the label value.

For the models described in this paper, we labeled approximately 4,000 RxMER samples. Severity of the impairments was not considered while building these proof-of-concept models.

## 4. Data Preprocessing


OFDM channels in Comcast typically vary in width from 48-96 MHz. This results in an approximate array of 940-1880 RxMER sub-carrier measurements per cable modem, per poll. Since the algorithms in consideration of this paper require a fixed input shape, the following options were evaluated:

- Fixing the width of the spectra to 1,880 by padding the end point with the following options:
  - Zeros
  - The average MER value
  - The last MER value
- Fixing the width of the spectra to a set value (e.g. 900) and applying a smoothing function to transform the raw input to the fixed length input.

As will be seen in the results section, the option of fixed spectra width of 1,880 with average MER value performed the best in our experiments.


## 5. Model Evaluation

The model evaluation strategy uses the typical train/test split, with 80% of the approximately 4k labeled RxMER samples to be used for training, and 20% kept out for validation. Figure 3 shows the distribution of samples by impairment. About 1,400 of the sample are normal -- i.e. they contain no visible impairment -- while the rest are distributed over several categories. Figure 4 shows the distribution of number of impairments by sample. Approximately 60% of the samples have a single impairment.



**Figure 3. Number of samples by label (impairment type).**





**Figure 4. Distribution of number of impairments/samples**

Since an RxMER instance for an OFDM Channel can have multiple impairments, detecting impairments is a multi-label classification problem. Predictions for an instance is a set of labels, and therefore the prediction can be fully correct, partially correct or fully incorrect. This makes model evaluation more challenging than binary classification problems, where accuracy, precision, recall and receiver operating characteristic (ROC) curves are typically used as evaluation criteria. In addition to determining accuracy, precision, recall and ROC for individual classes, we will be using the following evaluation criteria:

1. Exact Match Ratio (subset accuracy) – This indicates the percentage of samples that have all their labels classified correctly, given by:

$$\text{Exact Match Ratio} = \frac{1}{n} \sum_{i=1}^n I(Y_i = Z_i)$$

2. Hamming Loss – This indicates the fraction of labels that are incorrectly predicted, given by:

$$\text{Hamming Loss} = \frac{1}{n} \sum_{i=1}^n \frac{\text{xor}(Y_i, Z_i)}{|L|}$$

In the formulas above,

- $n$  is the number of multi-label samples
- $Y_i$  is the ground truth
- $Z_i$  is the prediction
- $L$  is the number of labels

The following are the definitions of Accuracy, Precision, Recall and F1 Score that will be used to evaluate the predictions of individual classes by a model:

$$\text{Accuracy} = \frac{(TP + TN)}{(TP + TN + FP + FN)}$$

$$Precision = \frac{TP}{(TP + FP)}$$

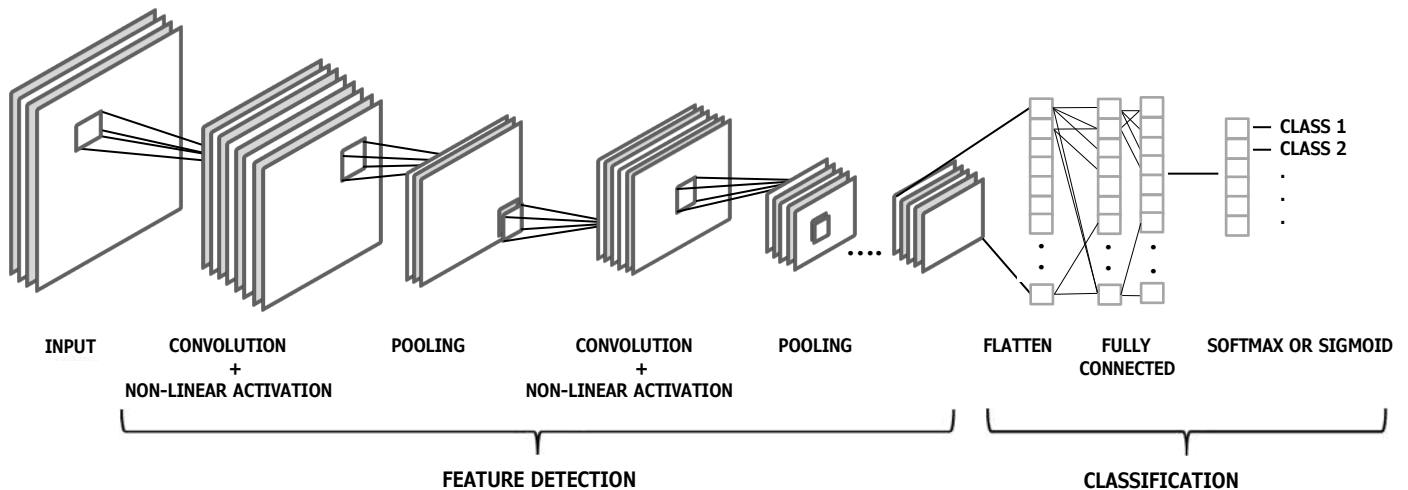
$$Recall = \frac{TP}{(TP + FN)}$$

$$F1\ Score = 2 * \frac{(Precision * Recall)}{(Precision + Recall)}$$

In the formulas above,

- True Positive (TP) – A true positive is an outcome where the model correctly predicts the positive class
- True Negative (TN) – A true negative is an outcome where the model correctly predicts the negative class
- False Positive (FP) – A false positive is an outcome where the model incorrectly predicts the positive class
- False Negative (FN) – A false negative is an outcome where the model incorrectly predicts the negative class.

## 6. Convolutional Neural Networks(CNNs)



**Figure 5. Convolution Neural Network (CNN) components**

CNNs are a class of neural networks that have proven to be extremely effective in recent years in the field of perceptual problems, specifically image recognition. They can be used to process data that has a spatial structure. In addition to being used for images that have a 2-D grid, they can also be used for 1-D structures, such as time series. In our use case, we use CNNs to classify impairments based on the 1-D RxMER per subcarrier array for OFDM Channels.

The model architecture is shown in Figure 5. CNNs contain the following components in addition to the input and output layers:

- **Convolutional layers:** Convolution layers use filters that perform convolution operations to scan the input layers with respect to dimensions. Convolution layers perform several convolutions in parallel, to produce a set of linear activations. Then, each linear activation is run through a nonlinear activation function, such as the rectified linear activation function, to generate a feature or activation map.
- **Pooling layers:** Pooling layers are typically applied after a convolution layer and perform a down sampling operation. They replace the output of the feature map with a summary statistic, such as max or average of the nearby outputs.
- **Fully connected layers:** Fully connected layers are typically present toward the end of a CNN architecture, and operate on a flattened input where each input is connected to all neurons.

There is no single network architecture, and understanding the network architecture is an area of great study. CNNs typically have a series of convolutional and pooling layers that are stacked together. The features that the convolutional/pooling layers detect increase in complexity as we go further down the network[1]. The convolutional and pooling operations provide translation invariance that uniquely distinguishes CNNs from other types of neural networks. Invariance refers to being able to recognize an object even when its appearance varies in some way.

## 7. Modeling

We experimented with CNN architectures that had between 1-3 convolution blocks and 1-3 fully connected layers. A grid search was performed on the following hyper-parameters before selecting the final model:

**Table 1. Hyperparameters, Ranges evaluated during training**

| Hyperparameter                            | Range                                                     | Hyper-parameter Type |
|-------------------------------------------|-----------------------------------------------------------|----------------------|
| Number of filters in convolutional layers | [32, 64, 96, 128]                                         | Network Structure    |
| Kernel size in convolutional layers       | [3,5,7,9]                                                 | Network Structure    |
| Pooling Size                              | [2,3,4]                                                   | Network Structure    |
| Fully connected hidden layer size         | [100, 150, 200, 250]                                      | Network Structure    |
| Dropout                                   | [0.3, 0.4, 0.5]                                           | Network Structure    |
| L2 Regularization                         | [0, 0.0001, 0.0005, 0.001, 0.005, 0.01]                   | Network Structure    |
| Learning Rate                             | [0.0001, 0.0003, 0.0005, 0.001, 0.003, 0.005, 0.01, 0.03] | Network Training     |
| Batch Size                                | [16, 32, 64, 128]                                         | Network Training     |

In addition, we evaluated the different methods to get an input of fixed shaped described in the data pre-processing section.


## 8. Results

We compared our models to 2 dummy models – one that has all labels assigned the most frequent class and another that randomly assigns labels based on the distribution in the training data. These 2 data points serve as benchmarks to compare against. In addition, we also compared the performance to a 3-layer conventional neural network.

**Table 2. Subset Accuracy, Hamming Loss for models**

| Model                                                      | Subset Accuracy | Hamming Loss |
|------------------------------------------------------------|-----------------|--------------|
| Model that predicts most frequent class                    | 0.319           | 0.163        |
| Model that randomly assigns labels based on distribution   | 0.146           | 0.202        |
| Neural Network – 3 dense layers                            | 0.645           | 0.047        |
| CNN – Zero padding at end to get 1880 Subcarriers          | 0.778           | 0.031        |
| CNN – Average MER padding at end to get 1880 Subcarriers   | 0.852           | 0.018        |
| CNN – Last MER seen padding at end to get 1880 Subcarriers | 0.824           | 0.022        |
| CNN – Fixed width of 900 Subcarriers                       | 0.809           | 0.024        |

As seen above, the CNN-based models significantly outperform the dummy and the neural network-based models on the validation dataset. While all the CNN based models had results in the same range, the model with average MER padding performed the best in our experiments.




**Figure 6. Network Architecture CNN model that had best performance on validation dataset [2]**

All the classes have accuracy above 90% on the validation dataset. Apart from amplitude ripple, they also have very good precision, recall and F1 scores. Due to its fine-grained nature, amplitude ripple does not perform as well and probably needs more training samples to get better scores. Another observation from the results is that incorrect classifications were often attributable to multiple impairments being present. This also indicates the need for additional labeled data.

Analog TV distortion and adjacent channel interference had fewer than 5 samples in the validation dataset and were not evaluated.

**Table 3. CNN model with average MER padding - Individual classes performance metrics on validation data**

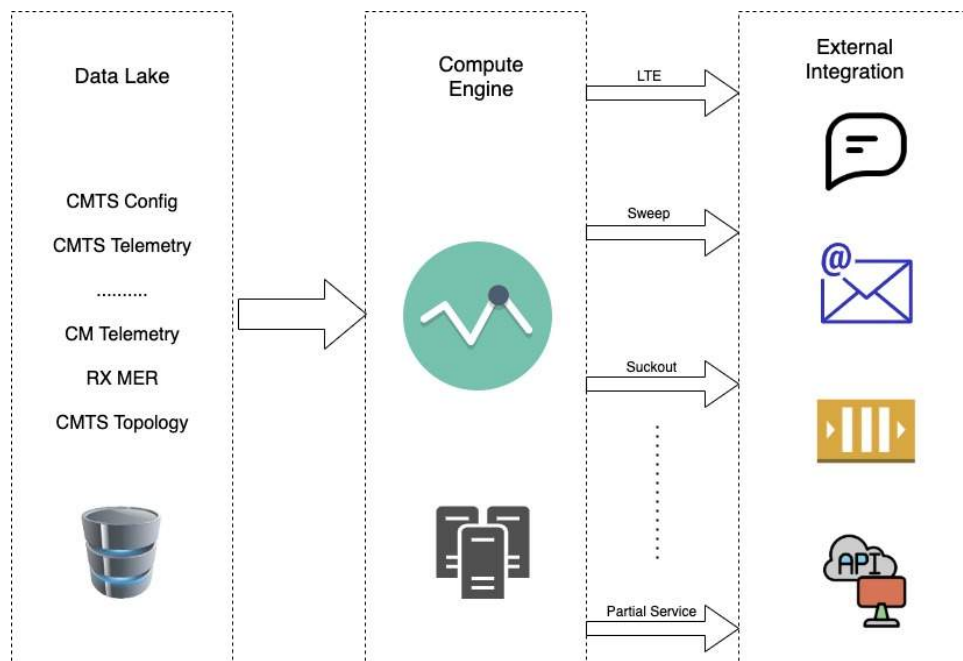
| Impairment             | Accuracy | Precision | Recall | F1 Score |
|------------------------|----------|-----------|--------|----------|
| Normal – No Impairment | 0.989    | 0.976     | 0.992  | 0.984    |
| Ingress                | 0.980    | 0.975     | 0.935  | 0.954    |
| Sweep                  | 0.997    | 0.933     | 0.933  | 0.933    |
| Roll-off               | 0.996    | 0.938     | 0.968  | 0.952    |
| Suck out               | 0.974    | 0.939     | 0.911  | 0.925    |
| Amplitude Ripple       | 0.939    | 0.912     | 0.715  | 0.802    |
| Low MER                | 0.974    | 0.958     | 0.924  | 0.940    |



**Figure 7. ROC Curves**

## 9. Machine Learning Pipeline

At a high level, the platform driving proactive network maintenance consists of 3 layers – Data Lake, Compute Engine and Integration Layer, as depicted in Figure 8.



**Figure 8. Machine Learning Pipeline Layers**

### 9.1. Data Lake

Comcast’s cloud-based data lake storage solution plays a pivotal role in our efforts to apply data science and develop analytical solutions to better understand and build a highly dynamic, resilient access network. The data lake also acts as a source of truth for a wide variety of data streams across Comcast. While the primary application of this data lake is to drive ML and analytical applications, custom integration and other solutions built on top of this data lake help power a wide variety of use cases. The platform is highly elastic, reliable and offers a rich set of tools for our analysts, engineers and data scientists to easily access this data and form a unified view of our network, customers and devices. It helps them use this data for analysis, visualization and ML applications, without the technical barrier of knowing the underlying infrastructure. The queries on the underlying data run on a Spark-based distributed computation engine, which is purpose-built to handle large data sources.

### 9.2. Compute Engine

The compute engine refers to the actual implementation of feature pipeline and ML models on polling data from the data collection framework. The standard ML life cycle involves model development, experimentation, training, test, validation and supporting model evolution. These ML models could be custom implementations developed from scratch or based on popular ML frameworks. During the development phase, the Compute Engine provides support to track and compare metrics for various experiments involving different models, features, tuning hyperparameters, etc. It also helps track code, data and model lineage while supporting promotion of models between different stages. To build the Compute Engine, we use open source ML frameworks and enhance it with custom functionality to suit our pipeline.

### 9.3. Integration layer:

The integration layer provides support to successfully integrate the recommendations from the pipeline to various OSS tools or directly to SMEs and technicians, through a broad set of interfaces and tools. The more seamless programmatic integration involves APIs and streaming events that would integrate with other OSS tools within Comcast. In severe cases, notifications could also be evented through Comcast's IM application, SMS and emails so they can be attended with greater urgency.

As a standard route, consumers would subscribe to the streams through Comcast's streaming data platform, which would event out our impairment notifications at pre-defined intervals. Once these events are consumed, they result in tickets/work orders being created automatically. This event-driven architecture enables any number of tools to easily integrate with our ML platform and track network events in real time. The following data are provided as part of notifications to assist with event prioritization and triangulation:

- Interfaces details such as PLC Location, Start/End Frequency, Total/Impacted cable modem counts
- List of impacted nodes
- List of impacted cable modems with severity of impairments
- Interface impairment rankings at the national, divisional and regional levels
- Reference to the API for historical data stored in the data lake related to the event
- Reference for API to enable fix agents to collect real-time on demand data, to confirm the issue is still present and to assist in isolation and confirm mitigation
- Other data sources to enrich the event, such as the mobile wireless carriers that overlap with the OFDM Channel

We also provide a real-time API that OSS tools integrate with and can be invoked on an on-demand basis. Once the impairments are attended to and a fix is identified, the network technicians can invoke the API through the UI of the OSS tools. Once a request is made to our API, the devices on the relevant interfaces/nodes are polled in real-time. Those polling results analyzed and scored through the model and a response is sent back to the UI, which indicates if the identified fix resulted in clearing the impairment. If the impairment is no longer seen, the OSS tool automatically clears the ticket. It is important to note that the compute engine handles model management for both offline scoring of the models for the entire footprint, and real-time scoring based on live polling data.

## 10. Conclusion/Next Steps

We've seen very promising results in being able to classify impairments for OFDM channels using CNNs on our validation dataset, with a high degree of accuracy. However, we need to continue to iterate and use additional labeled samples to address the following:

- The samples we've used may not cover the gamut of RxMER curves for the entire footprint
- There are some impairments that had very few examples
- Impairments that were incorrectly classified were often due to multiple impairments being present
- Detecting severity in addition to the impairment

A crowdsourcing approach to labeling that involves field technicians and other SMEs would be beneficial and increase confidence in the models being developed. SMEs are also needed for addressing edge cases, where the impairments were not always well-defined.

In the models/data pre-processing phase, there are a couple of areas that can be explored further:

- Increase the number of training samples by using data augmentation techniques such as flipping the RxMER curves horizontally or scaling the entire sample. We'll need to consider that some augmentations cannot be applied to all impairments, as it would modify the RxMER curve in a way that makes the impairment no longer applicable. For example, samples with roll-offs cannot be flipped horizontally.
- Impact of data pre-processing techniques on models needs to be understood better. For example, a sample having roll-off may be incorrectly classified as having suckout, due to the average MER padding added to get to the standard width.

We will also be extending this effort beyond OFDM channels to cover upstream, video and D3.0 channels. For OFDM channels specifically, building models for classifying impairments that include both RxMER and RxPower will be explored.

## Abbreviations

|       |                                               |
|-------|-----------------------------------------------|
| API   | Application Programming Interface             |
| CMTS  | Cable Modem Termination System                |
| CNN   | Convolutional Neural Network                  |
| D3.0  | DOCSIS 3.0                                    |
| D3.1  | DOCSIS 3.1                                    |
| FN    | False negative                                |
| FP    | False positive                                |
| HFC   | Hybrid Fiber-Coaxial                          |
| MER   | Modulation Error Rate                         |
| ML    | Machine learning                              |
| MTTR  | Mean Time To Repair                           |
| OFDM  | Orthogonal Frequency Division Multiplexing    |
| OFDMA | Orthogonal Frequency Division Multiple Access |
| OSS   | Operational Support Systems                   |
| PMA   | Profile Management Application                |
| PNM   | Proactive Network Maintenance                 |
| RxMER | Receive Modulation Error Rate                 |
| ROC   | Receiver Operating Characteristic             |
| QAM   | Quadrature Amplitude Modulation               |
| SCTE  | Society of Cable Telecommunications Engineers |
| SME   | Subject Matter Expert                         |
| SNR   | Signal-to-noise ratio                         |
| TN    | True Negative                                 |
| TP    | True Positive                                 |
| UI    | User Interface                                |



## Bibliography & References

1. *Visualizing and Understanding Convolutional Networks* - <https://arxiv.org/pdf/1311.2901.pdf>
2. *Net2Vis -- A Visual Grammar for Automatically Generating Publication-Ready CNN Architecture Visualizations*, Alex Bäuerle, Christian van Onzenoodt, Timo Ropinski  
<https://arxiv.org/abs/1902.04394>

# **DAA Field Deployment, Path to Scaling, and Digital Node Use Cases**

A Technical Paper prepared for SCTE•ISBE by

**Jorge Salinger**

VP, Access Architecture  
Comcast Cable Communications  
1701 JFK Blvd – Philadelphia, PA 19103  
+1 (215) 439-1721  
jorge\_salinger@cable.comcast.com

**Steve Sigman**

Director, Access Architecture  
Comcast Cable Communications  
1701 JFK Blvd – Philadelphia, PA 19103  
+1 (609) 685-3480  
steve\_sigman@cable.comcast.com

# Table of Contents

| Title                                                                    | Page Number |
|--------------------------------------------------------------------------|-------------|
| 1. Introduction.....                                                     | 4           |
| 2. Cable Network Evolution.....                                          | 5           |
| 2.1. Typical HFC Networks Today .....                                    | 5           |
| 2.2. Growth Projections.....                                             | 6           |
| 2.3. The Advent of DOCSIS 3.1 .....                                      | 8           |
| 2.4. The Analog Modulated Forward Link in HFC Networks.....              | 10          |
| 2.5. Description of Options for Digital Forward Link.....                | 10          |
| 2.6. Comparison of Options for Digital Forward Link .....                | 12          |
| 2.6.1. Option 1: RF remains in the headend.....                          | 12          |
| 2.6.2. Option 2: Digital-to-analog conversion is moved to the node ..... | 13          |
| 2.6.3. Option 3: Lower PHY is moved to the node.....                     | 13          |
| 2.6.4. Option 4: Entire PHY is moved to the node .....                   | 14          |
| 2.6.5. Option 5: Move PHY and MAC to the node .....                      | 14          |
| 3. Benefits of Distributed Architectures .....                           | 15          |
| 3.1. Improved performance .....                                          | 15          |
| 3.2. Increased Headend Equipment Density .....                           | 16          |
| 3.3. Integration of HFC and Fiber Services.....                          | 17          |
| 3.4. Migration Strategy .....                                            | 18          |
| 3.5. From Today to Virtual CCAP.....                                     | 20          |
| 4. DAA Components, Use Cases and Generations .....                       | 21          |
| 4.1. DAA Components .....                                                | 21          |
| 4.1.1. Advantages and disadvantages of discrete components .....         | 22          |
| 4.1.2. Key DAA discrete components .....                                 | 22          |
| 4.2. Key aspects of DAA interoperability.....                            | 23          |
| 4.3. Use Cases.....                                                      | 24          |
| 4.3.1. Outside plant.....                                                | 24          |
| 4.3.2. Inside plant.....                                                 | 25          |
| 4.4. Generational considerations .....                                   | 26          |
| 5. Conclusion.....                                                       | 26          |
| Abbreviations .....                                                      | 27          |

## List of Figures

| Title                                                                            | Page Number |
|----------------------------------------------------------------------------------|-------------|
| Figure 1: Examples of HSD service tier capacity increase over time .....         | 6           |
| Figure 2: Example of narrowcast service growth over time .....                   | 7           |
| Figure 3: Example of downstream SNR for a large population of cable modems ..... | 9           |
| Figure 4: Digital Forward - High-level Architecture .....                        | 11          |
| Figure 5: Block diagram for Option 1 .....                                       | 12          |
| Figure 6: Block diagram for Option 2 .....                                       | 13          |
| Figure 7: Block diagram for Option 3 .....                                       | 13          |
| Figure 8: Block diagram for Option 4 .....                                       | 14          |
| Figure 9: Block diagram for Option 5 .....                                       | 14          |
| Figure 10: Reuse of broadcast capacity across multiple RPNs.....                 | 16          |
| Figure 11: Single traditional HFC node.....                                      | 19          |

|                                                             |    |
|-------------------------------------------------------------|----|
| Figure 12: RPN deployment step 1 .....                      | 19 |
| Figure 13: RPN deployment step 2 .....                      | 19 |
| Figure 14: RPN deployment step 3 .....                      | 20 |
| Figure 15: RPN deployment step 4 .....                      | 20 |
| Figure 16: RPN deployment step 5 .....                      | 20 |
| Figure 17: DAA implementation components .....              | 21 |
| Figure 18: Functional CMTS-RPD inteoperability matrix ..... | 24 |

## List of Tables

| <b>Title</b>                                                                 | <b>Page Number</b> |
|------------------------------------------------------------------------------|--------------------|
| Table 1: SNR required for DOCSIS 3.1 .....                                   | 9                  |
| Table 2: Categories of options for implementing a digital forward link ..... | 12                 |
| Table 3: Distributed architecture headend density gain .....                 | 17                 |

# 1. Introduction

Cable operators have been actively converging video and data services into a common Converged Cable Access Platform (CCAP) platform for a few years now. This trend, which requires an evolution towards newer, more modern, and denser equipment, should free up space in the headend.

However, as the success of high-speed data and video-on-demand services continues its seemingly eternal growth, the evolution of the access network progresses relentlessly towards expanded capacity and ever-smaller service groups. As a result, the spectrum allocated to narrowcast services increases, driving operators to deploy Data Over Cable Service Interface Specification (DOCSIS<sup>®</sup>) services including 32 SC-QAM (Single Carrier Quadrature Amplitude Modulation) channels and at least a fraction of an OFDM (Orthogonal Frequency Division Multiplexing) channel. In some cases Cable operators have even gone farther, to 36 and even 40 SC-QAM channels, and moving beyond a single OFDM channel onto a second one. In addition, if free spectrum in the network is not available, operators support capacity growth by segmenting service groups into smaller and smaller areas.

These expansion trends result in a continuous growth of headend equipment, which is already starting to exceed the capacity that headend facilities can support.

Therefore, the above trends are now intractably linked to two additional evolutions: distribution of components of the access network and virtualization of the core network functions. Furthermore, the implementation of a Distributed Access Architecture (DAA) is different in places where the plant is being upgraded, such as where a migration to a deeper use of fiber or N+0 is being implemented, versus in Cable networks where the existing plant needs to be segmented or expanded.

This paper will begin by outlining the evolution of service provider networks, and then describe why and how the migration to a distributed architecture is necessary and beneficial. The paper will then expand into features that can be implemented with Distributed Architectures and discuss the topic of Virtualization. Finally, the paper will explore how the implementation of DOCSIS 4.0 could be implemented in Distributed Architecture networks.

The paper is divided into the following sections:

- Cable network evolution, including:
  - Typical HFC networks today and growth projections
  - Advent of DOCSIS 3.1
  - The analog modulated forward link, options for its implementation, and comparison of the various options
- Why and how distributed architectures are useful, covering the following key benefits:
  - HFC network performance improvements by migrating to digital transport, especially to maximize the use of the higher order modulation rates that DOCSIS 3.1 and beyond offers
  - Headend density increases, which is becoming critical as service providers are segmenting service groups more and more, including extending fiber deeper into the plant and implementing passive networks
  - Trunk fiber savings as we move to higher capacity digital links that can be muxed much more than analog links, and
  - The ability to eventually virtualize the remaining upper layers of the CCAP
- Describe a Distributed Access Architecture, including:
  - DAA components
  - Key aspects of DAA interoperability
  - Use Cases
  - Generations and the advent of DOCSIS 4.0

## 2. Cable Network Evolution

### 2.1. Typical HFC Networks Today

Most MSO's hybrid fiber-coax (HFC) networks have been designed with an upper spectral boundary of 750 or 860 MHz, while some are designed to support 1 GHz and other newer networks designed to support 1.2 GHz. For the more abundant 750 or 860 MHz networks, if not already fully utilized, it is expected that use of their capacity will be increased to the point of exhaustion. This will happen as a result of 1) increased DOCSIS® usage for even faster high-speed data (HSD) service tiers; 2) additional high-definition (HD) programs (for broadcast [BC] and especially narrowcast [NC] services, such as video on demand [VOD] and switched digital video [SDV]); and 3) new service additions such as internet protocol (IP) video and cloud-based digital video recorder (cDVR.)

In recent years the growth in, and demand for, HD programming has resulted in the need for allocation of large numbers of EIA (Electronic Industries Association) channels for HD services, both for BC (Broadcast) and NC (Narrowcast), which has filled every available portion of the spectrum. This is especially true for BC, where large numbers of programs are offered in HD format, while simultaneously the need for distributing the standard definition (SD) version has persisted. This has resulted in the need for use of 3x to 5x the number of EIA channels than previously required. For example, a typical digital multiplex including 10 to 15 programs would require an additional 3 to 5 EIA channels for the HD equivalent streams, even assuming the newer, more sophisticated multiplexing schemes available in the market. Of course not every program is available, or still sought by subscribers, in HD format. But very large numbers of them are, including 100 to 150 BC programs.

The above is also applicable to a great extent in systems utilizing SDV technology for content distribution. The difference is that the HD and SD versions of the program are not distributed unless a subscriber is requesting them, which reduces the marginal increase in capacity. Assuming that all programs are distributed in only one format, which is certainly a valid expectation for programs of low viewership, then the increase in capacity for a conversion from SD to HD would just be the increase in capacity required for the transmission of the HD program without requiring the simultaneous use of bandwidth for both formats.

Additionally, considerable spectrum is needed to deploy high-capacity narrowcast legacy video services, especially cDVR, and a full-array of HD video-on-demand services. For the former, initial observations suggest that network requirements for cDVR may be as high as 4x to 5x that of VOD, and that peak utilization overlaps, at least partially, with that of peak use for other narrowcast services.

Finally, the growth in HSD services continues. All network operators have offered increased service tiers and observed an increased use of HSD service capacity for well over a decade now, as shown in Figure 1, which amounts to a constant year-over-year compounded growth. The applications have changed throughout this time, and the demand has continued to increase at the same relentless rate.




Figure 1: Examples of HSD service tier capacity increase over time

How does this compare to other operator's data services and a longer period? Projecting an operator's HSD service growth back in time to when Internet services started 25+ years ago, services should have been about 100 bps. This coincides with the history of telephone modems from 110 and 300 baud modems from the mid-80s, to 56 Kbps/V.42, into ISDN (Integrated Services Integrated Network) services.

This demonstrates that the growth seen in MSO's HSD services is typical over a much longer period of time, rather than an exception observed by operators in recent years.

## 2.2. Growth Projections

From all of the above, it follows that, should the usage growth pattern continue at the same rate as in the past, networks will be required to provide >1 Gbps HSD services within the next few years. This growth, coupled with the surge in HD video formats, and more personalized narrowcast services, will result in a significant growth in narrowcast capacity, as shown in Figure 2.

To support this growth, MSOs have deployed bandwidth reclamation tools such as SDV for digital broadcast, digital terminal adapters (DTAs) for analog service reclamation, or a combination of both. These tools have been extremely valuable to MSOs, and their operational complexity and cost well justified.

In the case of SDV, early predictions from industry analysts projected that the efficiency of SDV would reach 40% (e.g., programs requiring 10 EIA channels could be carried in 6). This has proven to be understated, since it was based on the use of SDV to reduce bandwidth required for existing services. As SDV's role in the network grew, the efficiencies have been even greater, especially as SDV expanded in scope to support niche service introductions that have low initial viewership and would have otherwise been difficult to deploy.

The benefit of DTAs has been just as, or perhaps even more striking. MSOs deploying DTA devices are able to eliminate the need to distribute the analog channels in the network. Even if DTAs are

distributed to top tier analog customers, such as only the traditional expanded basic subscribers, the move would reduce a channel line up from perhaps 50 EIA channels dedicated to 50 analog programs to perhaps as little as 4 EIA channels dedicated to transport the 50 programs in their digital-equivalent transport. Using the same comparison method as the above SDV case, this is a >90% efficiency. If extended to the entire analog tier the efficiency gains are very significant.

Despite the availability of these tools, they are not universally applicable. With respect to SDV, in general it is not likely that all broadcast programs will be switched since experience shows that many broadcast programs are constantly viewed by someone in the service group during peak hours, which will leave a large portion of the spectrum still used for broadcast. Similarly, not all analog channels can be removed in the short term due to operational and/or cost constraints.

Additionally, while many MSOs will use one or both tools, in general these tools won't be used by every MSO for all applications.

Finally, there are also significant potential gains to be achieved from the use of advanced video compression standards (VCS), and variable bit-rate (VBR) multiplexing. In the case of VCS, coding efficiencies of approximately 50%, depending on implementation and content type, can be obtained with H.264 / MPEG-4 Part 10 . Furthermore, with the release of the H.265 standard in April of 2013, it is possible to achieve a 50% improvement over H.264. The use of VBR promises to result in a capacity efficiency gain of as much as 70% versus constant bit rate (CBR) . The combined gains from using the above approaches for multiple services are even more significant.




Figure 2: Example of narrowcast service growth over time

However, these are difficult tools to take advantage of, from a network perspective, because on a proportional basis, relatively few legacy set-tops will support all these technical advances, especially



H.265. These tools are more likely to find significant support in equipment designed to handle newer, IP-video based services.

And, this approach will nonetheless require additional capacity from the network. This is especially true when considering that the deployment of these advanced video services will result in an additional simulcast of video programs, at least initially. This is because, realistically, advanced services will not, at least initially, replace the currently deployed service formats

Furthermore, ubiquitous support for such devices would require considerable spectrum if the legacy services are maintained for an extended period-- which is expected, given that legacy devices will continue to be deployed, for some amount of time.. Moreover, this increase in simultaneous use of the more advanced IP video services while maintaining legacy services will be especially impactful over time, as the number of IP video services increases.

All of the above, coupled with the success experienced by MSOs in recent years with business services, homes security, etc., will likely require the deployment of IP capacity beyond what can be supported today. As well, it will require the development of tools for increased spectral efficiency and/or unleashing additional spectrum in the HFC network. The following sections of this paper will enumerate some ways in which this can be achieved.

### **2.3. The Advent of DOCSIS 3.1**

As it has been pretty well advertised in the media, DOCSIS 3.1 development has been quite extensive. Most MSOs announced deployments of DOCSIS 3.1 across their markets, and several operators have even deployed DOCSIS 3.1 throughout their entire footprint.

The key motivation for the 3.1 version of the DOCSIS specification is, in a nutshell, to scale DOCSIS more efficiently, both from the cost and operations perspectives.

While for the first 10 years of DOCSIS deployments it was possible to offer Internet services and support its growth with just one downstream DOCSIS channel, over the last 5-10 years services have required many more channels. This is because the year-over-year growth drove service speeds well above the initial levels, to speeds of 50 Mbps, 100 Mbps, and even much higher Mbps tiers today, which can't be supported by the single channel. Therefore, MSOs deployed multiple DOCSIS channels using DOCSIS 3.0, sometimes using 32 or more channels, and even requiring capacity beyond that supported by DOCSIS 3.0.

To that end, the 3 key goals and features of DOCSIS 3.1 are:

1. Much more efficient use of spectrum, with up to 50% improvement in bandwidth efficiency (or bps/Hz), resulting from:
  - a. The use of more efficient forward error correction (i.e., replacing the older and less efficient Reed-Solomon approach with the far more efficient Low Density Parity Check (LDPC), and
  - b. The addition of the higher-order modulations 1024 and 4096 QAM downstream and 256 and 1024 QAM upstream.

These new modulation schemes provide 2 and 4 bits per Hertz/second of improvement in both the upstream and downstream signal directions, while the use of the new forward error correction approach provides approximately 5 dB better RF performance. The end result is that MSOs are

able to transport 1 Gbps of DOCSIS capacity in about 120 MHz of spectrum. For context, doing the same with DOCSIS 3.0, using single-carrier QAM requires about 180 MHz of spectrum.

2. Cost reduction, mainly by leveraging technologies commonly used in other transmission media, such as the inclusion of Orthogonal Frequency Division Multiplexing, which is used extensively in wireless and wireline transmission media. Specifically, the addition of OFDM for the downstream and OFDMA (Orthogonal Frequency Division Multiple Access) for the upstream should enable MSOs to reduce costs by “packing” more bits in the HFC network more efficiently. As a result, these technologies will likely attract a larger supplier ecosystem, increasing innovation and fueling competition.
3. Enable a simple and orderly transition strategy, both with respect to compatibility with the previous generation of CMTS and CM equipment while simultaneously supporting an expanded spectrum capacity in the HFC network.

Specifically, DOCSIS 3.1 cable modems operate with DOCSIS 2.0 and 3.0 CMTS/CCAP equipment, enabling deployment of DOCSIS 3.1 CPE (Customer Premise Equipment) as soon as available. Similarly, DOCSIS 3.1 CCAPs support DOCSIS 2.0 and 3.0 CPE allowing MSOs to upgrade headend equipment without having to change any of the existing CPE. And, both DOCSIS 3.1 CM and CMTS equipment support the currently required upstream and downstream spectrum, plus an expansion of the upstream to 85 MHz and beyond, and of the downstream up to 1.2 GHz.




Figure 3: Example of downstream SNR for a large population of cable modems

Figure 3 depicts the downstream signal-to-noise ratio (SNR) as reported by a very large population of cable modems. This data shows that many cable modems will be able to support the high-order modulation profiles included in DOCSIS 3.1.

| Modulation            | Signal-to-Noise Ratio |
|-----------------------|-----------------------|
| 512 QAM               | 27 dB                 |
| 1024 QAM              | 30 dB                 |
| 2048 QAM              | 33 dB                 |
| 4096 QAM              | 36 dB                 |
| 8196 QAM <sup>1</sup> | 39 dB                 |
| 16384 QAM             | 42 dB                 |

Table 1: SNR required for DOCSIS 3.1

<sup>1</sup> 8196 QAM and 16384 QAM are included for future consideration in the DOCSIS 3.1 specifications

Assuming an 8/9 LDPC coding ratio, Table 1 shows the required SNR for the modulation rates included in DOCSIS 3.1:

Applying the SNR requirements from Table 1 to the population of modems shown in Figure 1, we can easily see that a large population of cable modems would not achieve sufficient SNR to operate at 4096 QAM. Furthermore, if sufficient headroom is allowed to account for environmental fluctuations, the population of cable modems that would not receive signals with sufficient SNR to operate at 4096 QAM would be significant.

## **2.4. The Analog Modulated Forward Link in HFC Networks**

As their name indicates, hybrid fiber-coax networks use a fiber transport between the headend and the coaxial cascade. This fiber link, intended to reduce the size of cascades, as a means to improve performance, was originally developed with analog modulated lasers and receivers in both signal directions, upstream and downstream.

Over time, the performance of the upstream link was improved by replacing the analog modulation with a digital transport. This change improved performance significantly, and allowed for longer distances between the headend and the node. Different vendors implemented their own methods and technical capabilities to implement a digital transport, which resulted in incompatible systems and required the use of the same vendors' components for both the node and the headend.

However, the downstream link remained almost unchanged over time, with the only enhancements focused on improving distance and RF spectrum capacity. Performance has not really been an issue like it was in the upstream.


But more importantly, while the digital capacity of the upstream was limited to a few Megabits per second, well under a gigabit of digital capacity which could easily be digitized and carried with Ethernet optics, the downstream digital capacity needed to transport the downstream spectrum has been considerably higher, reaching and even exceeding 10 gigabits per second.

Because of the above, analog forward links continue to be used. Even though headend equipment is currently capable of launching signals with >47 dB MER performance, which is sufficient to generate and transport 16,384 QAM signals, analog lasers are limited to about 35-38 dB of MER performance, which would limit end-of-line performance to barely enough for 2,048 QAM or 4,096 QAM in short cascades in the best of the cases.

## **2.5. Description of Options for Digital Forward Link**

As time has gone by, technology evolution and certain developments as described below have enabled options for implementing a digital forward link. These include:

1. Evolution of QAM edge modulators which have gone from single and/or a few modulators to supporting 32, 64 or even more modulators,
2. Development of the CCAP, combining the functions of the video QAM modulator and DOCSIS into a single platform, and
3. Migration to digital video, either partially or completely.



*Figure 4: Digital Forward - High-level Architecture*

With this technological evolution, it is conceivable to remove the RF combining network, and instead implement it digitally in the edge device, such as the CCAP.

This evolution of the edge headend devices makes it possible to envision several options for digitizing the forward link.

Fundamentally, the migration to a digital forward includes the components included in Figure 4, as follows:

- The headend device, such as a CCAP, which would be a high-density edge QAM comprising QAM modulation for the entire spectrum,
- The node would contain components normally implemented in the edge QAM or CCAP which generates the RF signals,
- The link between the headend device and the node would be comprised of a digital interface, such as an Ethernet link.

There are then various approaches for how a digital forward link can be implemented to replace the currently used analog link. These various approaches for distributing the various components can be categorized into 4 groups, plus 1 option that would still leave an RF generation at the headend device, as outlined in Table 2:

| Option                                                       | Description and Approach                                                                                                                                                                   |
|--------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1. Maintain RF output in the headend                         | <p>1.a Headend equipment remains unchanged</p> <p>1.b Headend RF output is digitized, transported digitally, and RF is regenerated in the node</p>                                         |
| 2. Remote the DAC from the PHY                               | <p>2.a The DAC is removed from the headend</p> <p>2.b Digital samples are transported digitally to the node where the DAC generates the RF signals</p>                                     |
| 3. Partition the PHY and remote the lower portion of the PHY | <p>3.a The PHY is split between the headend and the node</p> <p>3.b The digital bit stream between upper and lower PHY is transported from headend to node</p>                             |
| 4. Remote the entire PHY                                     | <p>4.a The entire PHY modulation is moved to the node</p> <p>4.b The MAC remains in the headend, and MAC frames are transmitted from the headend to modulator that resides in the node</p> |
| 5. Remote the entire PHY and MAC                             | <p>5.a The entire PHY and MAC is removed from the headend device and placed in the node</p> <p>5.b IP frames are transported from the headend to the node.</p>                             |

Table 2: Categories of options for implementing a digital forward link

## 2.6. Comparison of Options for Digital Forward Link

There are pros and cons for each of the options. The following sections outline these tradeoffs.

### 2.6.1. Option 1: RF remains in the headend




Figure 5: Block diagram for Option 1

- Equivalent to digital return, the RF output from the headend device is digitized, transported digitally, and converted back to RF in the node.
- Maintains HFC transparency
- This option results in the highest bitrate over fiber; the capacity for multiple nodes would not fit into the available capacity of one 10G fiber

- There is a loss of MER (Modulation Error Ratio) in the double conversion, so this option provides the least performance improvement
- Results in the least intelligence placed in the node, but an additional conversion stage is added in the headend

### 2.6.2. Option 2: Digital-to-analog conversion is moved to the node




Figure 6: Block diagram for Option 2

- Requires separation of the digital-to-analog conversion from the modulator
- Together with Option 1, results in the least intelligence in node
- Similar high bitrate over fiber as Option 1; capacity for multiple nodes would not fit into the available capacity of one 10G fiber

### 2.6.3. Option 3: Lower PHY is moved to the node




Figure 7: Block diagram for Option 3

- The PHY layer needs to be split into two components: upper and lower PHY
- More intelligence than in either of the previous options is placed in the node
- Although lower than the previous options, this option also results in a very high bitrate over fiber
- This option would require an industry proprietary point-to-point link between the headend port and the node to transport the I and Q samples
- Implementation of this option would require the definition of interfaces which have never been defined in previous versions of the DOCSIS specifications, which in turn would result in modification of the silicon used and/or planned to date, and therefore results in the highest implementation complexity

#### 2.6.4. Option 4: Entire PHY is moved to the node




Figure 8: Block diagram for Option 4

- More intelligence is placed in the node than with all previous options
- This option results in the lowest bitrate over fiber; multiple nodes fit into the capacity of a 10G fiber
- Enables a packet-based link between the headend and node, which results in significant benefits outlined later in this paper
- Could use existing/planned silicon devices, and thus may be the easiest and quickest to implement
- Offers the best MER performance improvement over analog

#### 2.6.5. Option 5: Move PHY and MAC to the node




Figure 9: Block diagram for Option 5

- This option puts the most intelligence in the node
- The data rate between the headend and the node is equivalent to the actual data transmitted, except for the addition of ancillary network data
- Same packet-based network benefits as Option 4
- Same highest MER performance as Option 4

Either of the 5 options described above accomplishes a migration away from the analog-modulated forward link and into a new era where the link between the headend and the node becomes a digital link. And, while either of the above approaches accomplishes a migration towards a digital link, over the years options 4 and 5 have received the most attention because of their relative implementation simplicity versus options 1, 2 and 3. We now call these options Remote PHY and Remote

MAC(Media Access Control)-PHY, and we call the devices that implement them RPDs (Remote PHY Devices) and RMDs (Remote MAC-PHY Devices) respectively.

As we migrated towards the implementation of a digital link, and separated either the physical layer in a Remote PHY implementation, or also migrated the MAC in a Remote MAC-PHY implementation, we stepped into the era of the Distributed Architectures. In these Distributed Architectures, the remainder of the CCAP in the headend no longer needs to be implemented in an application-specific hardware design. Instead, the remainder of the CCAP in the headend can be implemented entirely as software running in general purpose compute platforms, which we now call a Virtualized headend platform.

The next sections of this paper will focus on the benefits of a distributed architecture, discuss some of the features of distributed architectures, and outline network evolution strategies.

### **3. Benefits of Distributed Architectures**

There are many benefits from the implementation of Distributed Architectures. The following sections of this paper describe them.

#### **3.1. Improved performance**

Improvements on performance are achieved in multiple ways, including:

- Improved SNR characteristics
- Longer link distances
- Higher reliability
- Better use of capacity

As described in the above sections of this paper, one key benefit of Digital Forward Link is the improved performance resulting from the migration from an analog to a digital link. This gain varies depending on the characteristics of the analog link being replaced, but can be generalized as an improvement of 5 dB in signal-to-noise ratio at the end of the line. This gain will result in higher capacity/Hz as it will be possible to run the higher order modulations as shown in Figure 3 and Table 1 for more of the cable modems in the network. This will enable significantly higher transport capacity for customers in the HFC network.

In addition, the Digital Forward Link will enable longer distances between the headend and the node. This is because digital interfaces, such as an Ethernet link, are designed to operate over much longer distances while carrying the designated capacity. Extending the distance between the CCAP and the Digital Node would enable MSOs to move their CCAP devices back in the network to more centralized facilities, leaving the hub or OTN free of CCAP equipment. The benefit of such change could be very big for some MSOs, especially as segmentation of the network continues towards smaller service groups, for which additional CCAP equipment needs to be deployed.

A third benefit from the Digital Forward Link is improved reliability of the optical link. It is well known that analog links require periodic maintenance and are subject to the effects of environmental changes. By contrast, Ethernet optical links are far more stable across a wider range of environmental conditions, and require little to no maintenance. The impact of this benefit could be very significant to MSOs.



Finally, the data transmitted through the link can be used more efficiently. One key example of such efficiency is the case where one link is used for multiple remote devices. As shown in Figure 10, one link from the headend CCAP device can be used to transport broadcast services once for multiple remote devices. This is achieved by using multicast addressing, whereby each of the remote devices uses the same lineup for each of the respective service groups. In doing so, a single link from the headend CCAP can be used for all the remote devices without exhausting the transport link capacity.




Figure 10: Reuse of broadcast capacity across multiple RPNs

### 3.2. Increased Headend Equipment Density

The implementation of distributed architectures makes it possible to improve the density of CCAP devices in several ways.

First, while CCAP devices are normally implemented via separate upstream and downstream line cards, a distributed architecture line card could implement both upstream and downstream. This, in effect, doubles the capacity of a CCAP chassis.

In addition, a typical CCAP downstream line card will house 8 or perhaps 12 RF ports, as defined by the printed circuit board space consumed by the components required for RF modulation, plus the sheer connector spacing required. However, Ethernet connectors can be placed considerably closer to one another, allowing a similar line card to easily house 16 to 24 ports. This additional density gain once again doubles the capacity of a CCAP chassis.

Finally, it is possible to consider “daisy chaining” digital nodes (RPNs) off of a single CCAP Ethernet port. This is because, on the one hand, the capacity of an 10 Gbps Ethernet link would support the capacity needed for a single RPN. Plus, it is possible to generate an RPN “channel lineup” by transmitting the broadcast content once to multiple RPNs. As depicted in Figure 10, the data stream transmitted from the CCAP could contain a single “copy” of the broadcast line-up content, plus individual versions of the narrowcast content for each of the RPNs. The RPNs would then reuse the broadcast lineup content to recreate the individual RPN channel lineup. In this way, each service group served by the CCAP port would contain the same broadcast lineup while allowing for its own unique narrowcast line-up.

Then, as the narrowcast lineup capacity grows over time, CCAP ports would be segmented to support less RPNs, akin to the way service groups are split today to provide more narrowcast capacity as it is required.

As summarized in Table 3 below, the combined effect of the 3 factors described above is very significant, ranging from 8x to 18x of headend capacity gain. From a space and power perspective, this can facilitate huge savings.

| Density Factor                        | Density Gain |
|---------------------------------------|--------------|
| Combined US/DS line card              | 2x           |
| Greater number of ports per line card | 2x to 3x     |
| Multiple RPNs per CCAP port           | 2x to 3x     |
| Combined capacity gain                | 8x to 18x    |

*Table 3: Distributed architecture headend density gain*

But, just how meaningful is this headend density gain?

Consider: A migration from an HFC architecture with an average of N+5 (meaning an optical-to-RF node followed by 5 cascading amplifiers) to N+0 would require about 10x the number of nodes, and the headend density benefits resulting from the distributed architecture would neutralize the potential increase in CCAP equipment.

It is then quite clear that from a space and power savings, distributed architectures take the benefit of CCAP to a whole new level.

### **3.3. Integration of HFC and Fiber Services**

One of the largest areas of growth for MSOs is business services. MSOs have deployed business services via both cable modems and fiber-based infrastructure. The fiber-based services are either point-to-point, using Ethernet and wave-division multiplexing (WDM), or point-to-multipoint, using PON (Passive Optical Network) technologies (either EPON [Ethernet Passive Optical Network] or GPON [Gigabit Passive Optical Network]).

This duality results in the existence of two parallel networks. One of them, the HFC infrastructure, uses fiber from the headend to the node via an analog modulated link for the forward direction and either analog or proprietary digital return, followed by coax infrastructure from the node to the home. The other consists of digital fiber from the headend to the subscriber, which is used for commercial services.

Given the use of a digital fiber in both the forward and the return for the RPN, and especially because this digital fiber is based on Ethernet technology, it is possible to collapse both of these networks into a single infrastructure. Even without fully collapsing the Ethernet network better utilization of physical fiber can be enabled by the move to common DWDM wavelengths and multiplexers.

Therefore, the implementation of RPNs with an Ethernet interface between the CCAP and the RPN would make it possible to implement a PON interface at the RPN.

The benefits from this integration include:

- Reduction of the optical link for PON to the distance between the node to the customer premise

- The typical distance from a node to a customer premise in an N+0 architecture is 1-2 kilometers. This would virtually eliminate any distance limitations for PON, making it possible to implement the largest possible densities.
- In addition, this shortened distance would enable the use of lower power optics, which can translate into significant savings -- especially for 10 Gbps optics, and especially for the upstream, which results in significant savings in the ONU.
- Leverage a single network for multiple services, which will reduce maintenance and increase operational efficiencies.

### 3.4. Migration Strategy

Clearly, one of the more concerning issues to MSOs is the migration strategy.

Any migration that requires synchronized cut-overs, or which requires changes in multiple locations to execute, is problematic, and usually results in a barrier to adoption. Therefore, it is very important that the migration to distributed architectures allow for unsynchronized changes.

Furthermore, ideally the migration to distributed architectures allows for opportunistic changes in the network. For example, one such change would be to migrate a single node, such as would be the case in an MDU to increase capacity.

As it turns out, distributed architectures enable such gradual, unsynchronized and opportunistic changes in the network. What follows is an overview of the steps and components involved in the migration to distributed architectures.

Starting with the components of the network on both sides of the distributed architectures, neither the back-office nor the various components in the customer premise need to be modified in any way. All back-office components are unaffected by the migration to distributed architectures, and any additional MIBs for management and/or commands for configuration as needed can be added well before the first distributed architecture CCAP line card or node are deployed. With respect to customer premise devices, these would not be affected in any way in order to deploy distributed architectures, and any enhancements that are made possible through the introduction of distributed architectures would be implemented in CPE equipment that can be introduced before or after the migration to distributed architectures.

The critical portion of the network where changes need to be made are in the headend and the plant.

To begin with, the changes required in the headend are primarily in the CCAP platform. The CCAP architecture was specifically designed to support multiple technologies simultaneously, which makes it possible to install regular RF upstream and downstream line cards and distributed architecture line cards in the same chassis. While some MSOs may choose to deploy a separate CCAP platform for distributed architectures, it is certainly possible to support both types of line cards in the same chassis. Of course, these distributed architecture line cards can be installed at any time prior to beginning the migration in the plant, and any removal of RF upstream or downstream line cards can follow the deployment of any number of distributed architecture line cards or nodes.

Turning our attention to the plant, it is similarly possible to migrate regular nodes to distributed architecture nodes in any sequence. As an example, what follows is a sequence of steps where a single node is gradually converted from standard HFC to distributed architecture.

Figure 11 depicts a single HFC node connected to a CCAP device.




Figure 11: Single traditional HFC node

Figure 12 shows how the HFC node would be converted to RPN while the rest of the HFC network remains unchanged. The distributed architecture line card in the CCAP would have been deployed in the headend in a prior activity, and even the RPN could have been deployed before the day of the cutover. Then on the day of the change the fiber cable could be swung in the headend from one AM laser to the CCAP distributed architecture card, and in the field from the HFC node to the RPN. Of course it is not necessary to perform the migration in such a fashion, but it would be possible if desired.




Figure 12: RPN deployment step 1

Figure 13 depicts a possible step 2 in the process, whereby additional RPNs are installed to segment the original service group further. These additional RPNs could be daisy chained from the original RPN node by taking advantage of the broadcast reuse feature, minimizing complexity in the deployment process.

NOTE: The example depicted is one in which fiber is run to every amplifier station. However, a more efficient segmentation scheme would include optimal placement of RPNs in an N+0 HFC architecture with some turnaround of passive components.




Figure 13: RPN deployment step 2

Figure 14 shows how further segmentation could take place by replacing the remaining amplifiers in the network with RPNs.




Figure 14: RPN deployment step 3

Figure 15 shows that the RPN service group depicted above is segmented as additional narrowcast capacity is required. In this example, 2 of the RPNs from the distributed architecture service group shown in Figure 14 are split into separate service groups using separate CCAP ports.




Figure 15: RPN deployment step 4

Eventually each of the RPNs could be connected to an individual CCAP RPN port as shown in Figure 16. This would provide up to 10 Gbps of capacity to each RPN. This could, for example, be desirable to provide both RF and PON services from the RPN.

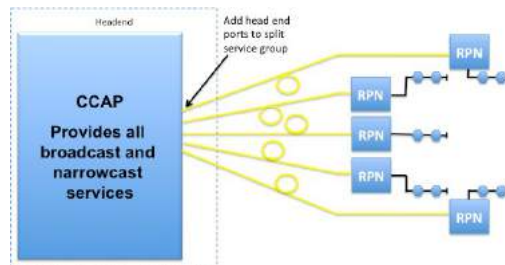


Figure 16: RPN deployment step 5

Similarly, the distributed architecture line card in the CCAP could be upgraded to support even more capacity as such capacity is needed and becomes cost effective. For example, the Ethernet link from the CCAP to the RPN could eventually be upgraded to 40 or 100 Gbps, both of which are already commercially available.

### 3.5. From Today to Virtual CCAP

As the network has to continue in operation through the transition, virtualizing the CCAP requires careful planning and a sensibly staged process. As with roads, where cars must be kept moving during any lengthy highway reconstruction, in the network customer traffic must continue flowing day after day. In a sense, while road work is visible to car drivers, in a network the modifications remain invisible to the end user.

One way to do so is to migrate individual functions, one at a time. So, one must develop a list of the functions that would be virtualized, and this list would be prioritized, such as on the basis of complexity of implementation and benefit. Those features with the lowest implementation complexity and the highest benefit would be prioritized higher in the list, and consequently implemented first.

In DOCSIS 3.1, one of the functions that would rise to the top of any such list is Modulation Profile Management (MPM). This is because MPM will take time to be implemented by vendors in a CCAP chassis, but implementing externally via virtualization could be quite simple. In the process, its benefit to operators is quite significant since it would enable better efficiencies from DOCSIS 3.1.

Over time, implementing virtualization of the various functions of the CCAP would lead to a completely virtualized CCAP platform. Such a platform would be more easily scalable than CCAP platforms are today, where segmentation of service groups requires the addition of more chassis in a linear relationship fashion.

In addition, and perhaps more importantly, virtualizing the CCAP will enable the development of additional functionality, and improvements to such functionality, to occur much more rapidly than it is possible to do today.

## 4. DAA Components, Use Cases and Generations

As MSOs move forward with the implementation of Distributed Access Architectures, there are already many useful lessons learned. The following sections of this paper describe three key aspects: DAA components, use cases, and the generational aspects of DAA.

### 4.1. DAA Components

As the link from the headend to the node is converted from analog modulated forward to digital, using Ethernet as the transport, several approaches can be taken for the implementation of the remaining headend components. In this section we will examine one approach in some detail, for which a key goal is to convert all required components into functionally individual software pieces implemented independently.

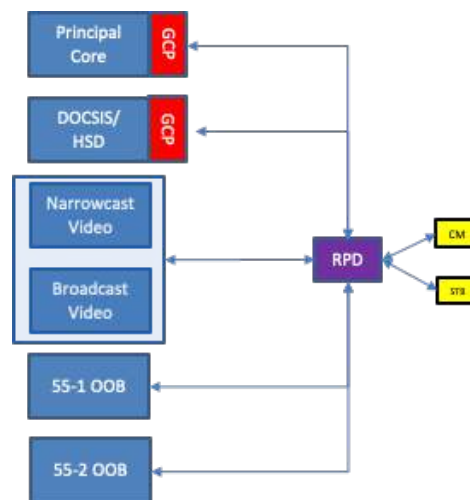


Figure 17: DAA implementation components

#### **4.1.1. Advantages and disadvantages of discrete components**

As shown in Figure 17, an implementation approach for DAA is to develop discrete SW components for each of the various DAA components. Some of the advantages for doing so include:

- The implementation can consist of a multi-vendor platform, where each component can be developed by a different party.
- By having smaller functional components their implementation tends to be simpler.
- Time to market also tends to be reduced.

However, implementation of smaller discrete components has its downsides, such as:

- There is an implied requirement to more tightly specify the behavior of each component to ensure that the overall system will operate as intended.
- Interface specifications between the various components is required.
- Management of the various components, including their configuration and upgrade is generally more complicated, and requires more elaborate orchestration.

#### **4.1.2. Key DAA discrete components**

The key components depicted in Figure 17 include:


- The Principal Core, which is the first component that the RPD will contact after receiving an IP address.
  - The Principal Core is implemented such that it will configure all the RPD functions except DOCSIS channels and behavior, which will be implemented by the DOCSIS CMTS.
  - As depicted in Figure 17, the Principal Core communicates with the RPD using the GCP protocol, for which it is known as the GCP Principal, or GCPP for short.
  - Included in the GCP Principal Core are all the non-DOCSIS command and control functions for the RPD, including configuration, management and reporting.
- The DOCSIS Core, which is the second component that the RPD will contact in the network.
  - The DOCSIS Core also communicates with the RPD using GCP.
  - The DOCSIS Core provides all configuration, command and control for DOCSIS channels, both downstream and upstream.
- Narrowcast and Broadcast Video engines

- Implemented as separate components, the Narrowcast and Broadcast Video engines provide all the content services for the various RPDs in the network.
- Neither the Narrowcast nor the Broadcast Video engines communicate with, nor have knowledge of, the RPDs.
- Services are configured statically in the Narrowcast and Broadcast Video engines upon their bring-up, and are multicasted to all RPDs, which listen for these services as configured by the GCP Principal.
- The Narrowcast and Broadcast Video engines could be implemented separately, but they could be operated together as a single functional system.
- Out-of-Band engines or cores
  - The OOB (Out Of Band) functions are implemented separately from the video engines.
  - Given that video systems are implemented using a single encryption and command/control technology, only one (i.e., either SCTE 55-1 or SCTE 55-2) of them is deployed in any one system.
  - The OOB function may or may not implement GCP for communicating with the RPD. When GCP is implemented the OOB server is a Core, and it will configure the OOB downstream and upstream OOB channels in the RPD. However, when GCP is not implemented the OOB server is an engine, and the GCP Principal will configure the downstream and upstream OOB channels.
- Finally, not depicted in Figure 17 is a very important component: the Timing Server.
  - The Timing Server, also known as the Grand Master, provides the critical timing synchronization for all the DAA components.
  - Each of the DAA components will include a Timing Client, which will communicate with the Timing Server to maintain timing synchronization.
  - While timing synchronization is not absolutely critical for video services, it is imperative for DOCSIS service to operate. Therefore, video services may be initiated before timing synchronization is achieved, but DOCSIS services will not.

## 4.2. Key aspects of DAA interoperability

The base implementation of a DAA system is generally simple. However, significant complexity is introduced when interoperability with different vendors' components are introduced.





*Figure 18: Functional CMTS-RPD interoperability matrix*

As depicted in Figure 18, the number of combinations of interoperable components increases geometrically as additional components are added on either side of the interoperability matrix. Having a single CMTS interoperate with multiple vendors' RPDs is complex and requires a lot of careful planning and implementation. If the number of CMTS implementations is increased to 2 or 3, the interoperability complexity doubles and triples respectively.

When considering the overall CCAP system, the complexity to achieve multi-vendor interoperability is even larger. For example, if multiple GCP Principals and/or multiple video engines and/or multiple OOB engines/cores are introduced into the mix, the amount of complexity and work required for testing and interoperability results in increases by orders of magnitude.

Therefore, a multi-vendor RPD deployment coupled with a single headend implementation is a sensible approach to an interoperable DAA ecosystem.

### **4.3. Use Cases**

In the same way as there are different kinds of nodes for different HFC network applications, there are Remote PHY devices with different characteristics that are best suited for each of the specific HFC network use cases. Similarly, while there are use cases for Remote PHY devices in the outside plant, there are also applications for Remote PHY devices in headends, or "inside plant" as it is frequently called, which will have different implementation characteristics. The following sections cover the key scenarios.

#### **4.3.1. Outside plant**

The environmental characteristics of Remote PHY devices developed for outside plant make the design of such devices very different than for RPDs developed for inside plant. The key characteristics for outside plant RPDs are as follows:

- Designs must conform to very tight space availability requirements inside of a node enclosure
- RPDs must support an environment where heat dissipation without the use of fans is critical
- Powered from quasi-square wave power supplies used in HFC networks
- Minimize power consumption to the extent possible given the limited amount of power during normal operation and especially during stand-by power mode

In addition, and perhaps more importantly, there are different kinds of nodes for different HFC network applications, which will drive varying designs for RPDs for outside plant, as follows:

- Traditional HFC networks include cascades of multiple amplifiers and cover a plant footprint of a few hundred homes. In such cases the network capacity offered by the Remote PHY device should be maximized, such as including multiple downstream and multiple upstream ports.
- Newer HFC networks are built with less, or even no amplifiers, and are targeted to cover smaller network footprints. In such cases it is not necessary for the RPD to support much more than a single downstream port, with either a single or dual upstream ports.
- Finally, given that scaling is needed as in any other network application, it should be possible to support greater capacity over time to the extent possible. For example, while initial deployment may only require a single downstream and/or upstream, over time service group segmentation may require additional downstream and/or upstream ports in a single node. For that purpose it is usually a design requirement that multiple individual RPDs fit into a single node enclosure.

Given the above, Remote PHY devices that are built with a single downstream and a single upstream (frequently called 1x1) or a single downstream with dual upstreams (1x2), such that MSOs can place a single one in the node or place additional units when capacity demands require it. Newer silicon designs include more capacity at lower power levels, making it possible to develop RPDs that contain multiple downstream ports and multiple upstream ports, such as 2x2 and 2x4 designs in a single Remote PHY device.

#### **4.3.2. *Inside plant***

In contrast to the outside plant environmental characteristics, Remote PHY devices developed for inside plant have other constraints that make the design of such devices very different than for RPDs developed for outside plant. The key differences in the design of RPDs for inside plant are as follows:

- Rather than the physical volume of the allocated space, the layout is a primary concern so that dense set-ups in a rack are possible, including cabling distribution in the front and/or back of the rack
- Designs must support an environment where forced air is used for heat dissipation, requiring airflow from front-to-back or back-to-front, sometimes allowing airflow from side-to-side and/or in a vertical direction
- Powering frequently requires DC power supplies, but AC power supplies are used in other cases

In the case of inside plant RPDs, these are frequently implemented in one of two different form factors:

- Modular, where RPDs are individually removable in a chassis-based design, or
- Fixed, where the entire set of RPDs are part of a monolithic device

The modular design is generally used in larger headends where the ability to replace a defective individual unit is a paramount concern. The fixed design is targeted for a smaller facility, or even a cabinet, where space is the primary concern.

#### **4.4. Generational considerations**

One final consideration for this paper is the evolution of the Digital Access Architecture to support new generations of equipment.

The initial implementation of the DAA components included support for DOCSIS 3.1. The main component that is specifically developed and implemented for DOCSIS 3.1 is the Remote PHY device, which incorporates ASIC devices which support up to DOCSIS 3.1, but will not support DOCSIS 4.0 functionality.

As newer parts of the DOCSIS 4.0 specifications are implemented, such as FDX, the RPDs will have to be swapped out in order to expand their support for DOCSIS 4.0. This process is akin to what had to be done with CCAP linecards in the past, where either upstream and/or downstream linecards are swapped over time as new versions become available. And, as part of this upgrade, the older equipment is reused in other locations where the newer equipment is not yet needed.

However, for the remaining DAA components, if these are implemented in software on general purpose compute platforms, these should be upgradeable to support newer DOCSIS specifications such as DOCSIS 4.0 and/or other enhancements by simply expanding the functionality implemented in software and downloading it to the platforms in which they run.

In fact, the process for upgrading the DAA components to support changes in the functionality including enhancements to DOCSIS, becomes easier than ever before given the nature of the DAA platform, especially when the DAA implementation includes a minimalistic Remote PHY device at the edge of the HFC network and virtualized components for the remainder of the DAA components.

## **5. Conclusions**

Demand for more narrowcast service capacity has driven many changes for MSOs to reclaim spectrum. Splitting nodes to reduce the number of homes passed in an HFC node's footprint is another method used to provide more narrowcast capacity. Enabling narrowcast services on this reclaimed and newly added spectrum requires more equipment in the headend. Given the trajectory of growth using existing equipment will create demand for more headend space. Deploying DAA allows for this growth without expanding facility footprint.

DOCSIS 3.1 created a demand for higher MER out of the node to take advantage of the highest modulation orders available. Transitioning to a digital forward link is the enabling function to achieve this demand. Moving the entire PHY layer to the node has become the standard method to implement a digital forward link.

Deploying DAA in the outside plant allows for more efficient infrastructure utilization between services. DAA and commercial services can share physical fiber due to common DWDM wavelengths and spacing. There is also the possibility of a converged Ethernet switching network to serve both DAA nodes and commercial customers. Having Ethernet inside of a DAA node allows for efficient PON deployments sourced from the same node.

DAA migration scenarios were outlined showing how RPNs can be initially deployed and scaled to offer more bandwidth in their lifecycle. This includes capacity on both the RF and ethernet sides.

Moving from specialized CCAP hardware to a virtualized system can lead to a more scalable system where scale can be added without the step function of adding new CCAP chassis. Feature velocity can be improved on a virtualized system.

Deploying DAA can be done with a CCAP core that provides all needed services or with discrete components that allow for a multi-vendor platform that includes smaller functional components. Implementing discrete components requires well behaved components that adhere to a well-defined specification. Key DAA functions include a GCP principal core, a DOCSIS core, a narrowcast and broadcast video source, an SCTE 55-1 or 55-2 OOB source, and PTP timing distribution.

Remote PHY devices are available for both inside and outside plant applications. Power level and segmentation capability / capacity are key attributes of outside plant RPDs. Density and serviceability are key attributes of inside plant RPDs.

Even though current generation RPDs lack support for FDX operation the other DAA components may have the ability to be upgraded for FDX operation using the current hardware. This enables an easier transition to FDX on an as needed basis later by swapping hardware on desired nodes while utilizing the existing platform.

## Abbreviations

|      |                                         |
|------|-----------------------------------------|
| A/D  | analog-to-digital                       |
| AC   | alternating current                     |
| AM   | analog modulated                        |
| ASIC | application specific integrated circuit |
| BC   | Broadcast                               |

|        |                                                  |
|--------|--------------------------------------------------|
| CBR    | constant bitrate                                 |
| CCAP   | converged cable access platform                  |
| cDVR   | cloud digital video recorder                     |
| D/A    | digital-to-analog                                |
| DAA    | distributed access architecture                  |
| dB     | Decibel                                          |
| DC     | direct current                                   |
| DOCSIS | data over cable services interface specification |
| DTA    | digital terminal adapter                         |
| DS     | downstream                                       |
| EIA    | Electronics Industry Association                 |
| Gbps   | gigabit per second                               |
| GCP    | Generic Configuration Protocol                   |
| GHz    | gigahertz                                        |
| HD     | high definition                                  |
| HFC    | hybrid fiber-coax                                |
| HSD    | high speed data                                  |
| Hz     | Hertz                                            |
| IP     | Internet protocol                                |
| ISDN   | integrated services digital network              |
| Kbps   | kilobits per second                              |
| MER    | modulation error ratio                           |
| MSO    | multiple system operator                         |
| NC     | narrowcast                                       |
| OFDM   | orthogonal frequency division multiplexing       |
| ONU    | optical network unit                             |
| OOB    | out-of-band                                      |
| OTN    | optical termination node                         |
| PHY    | physical                                         |
| PON    | passive optical network                          |
| QAM    | quadrature amplitude modulation                  |
| RF     | radio frequency                                  |
| RPD    | remote PHY node                                  |
| RPN    | remote PHY nodes                                 |
| SC-QAM | single carrier QAM                               |
| SD     | standard definition                              |
| SDV    | switched digital video                           |
| SNR    | signal-to-noise ratio                            |
| US     | upstream                                         |
| VBR    | variable bitrate                                 |
| VCS    | video compression standards                      |
| VOD    | video on demand                                  |
| WDM    | wave division multiplexing                       |

# **A Method and Framework for IoT Network Security**

A Technical Paper prepared for SCTE•ISBE by

**Umamaheswar Achari Kakinada**

Principal Engineer, Wireless R&D  
Charter Communications, Inc  
6360 Fiddlers Green, Greenwood Village, CO 80111  
847-544-6560  
Achari.Kakinada@Charter.com

**Dr. Hossam H. Hmimy**

Sr. Director – Wireless R&D  
Charter Communications, Inc.  
6360 Fiddlers Green, Greenwood Village, CO 80111  
720-536-9396  
Hossam.Hmimy@Charter.com

**Manish Jindal**

GVP – Wireless R&D  
Charter Communications Inc.  
6360 Fiddlers Green, Greenwood Village, CO 80111  
303-793-4486  
Manish.Jindal@Charter.com

# Table of Contents

| Title                                                | Page Number |
|------------------------------------------------------|-------------|
| 1. Introduction .....                                | 3           |
| 2. Proposed IoT Security Framework .....             | 5           |
| 2.1. Connect:.....                                   | 6           |
| 2.2. Monitor: .....                                  | 6           |
| 2.3. Analyze: .....                                  | 7           |
| 2.4. Compliance: .....                               | 7           |
| 2.5. Improvement cycle: .....                        | 7           |
| 3. IoT Data : Security, Integrity and Ownership..... | 7           |
| 3.1. Security and Integrity: .....                   | 8           |
| 3.2. Ownership: .....                                | 10          |
| 4. Standards support.....                            | 11          |
| 5. Conclusion .....                                  | 11          |
| 6. Acknowledgements .....                            | 11          |
| Abbreviations.....                                   | 12          |
| Bibliography & References .....                      | 12          |

## List of Figures

| Title                                                                    | Page Number |
|--------------------------------------------------------------------------|-------------|
| Figure 1 - IoT Security Life Cycle .....                                 | 5           |
| Figure 2 - Chain of Custody for IoT Data Security .....                  | 9           |
| Figure 3 - Message Flow for Chain of Custody for IoT Data Security ..... | 10          |

# 1. Introduction

The Internet of Things (IoT) is connecting any device or “thing” to the Internet. These devices include sensors, actuators, machines and wearables that are capable of collecting and transmitting data about location, activity, situational awareness, environmental activity, etc., to allow autonomous or real-time changes that can enhance or optimize everyday activities.

The number of “things” connected to the Internet is growing at nearly an exponential rate and delivering data about everything from the temperature of an industrial refrigeration unit, to household water consumption, to your heart rate. Each one of these “things” is a doorway into a larger network. Together, these generate colossal amounts of data and can add enormous value to many spheres of society. To realize this value, these things and the systems they are connected to and the data that is generated must be reliable and trustworthy. Thus, it is of paramount importance to secure the things, associated systems and the data generated.

More and more connected devices are being deployed, and a threat or security breach in one area or in one device can have a domino effect on other devices connected to that same network. Potential damage can extend past the network and the data. Finances, reputation, brand and disruption of city services are just some of the areas to be considered in a security assessment to mitigate potential risk. We’ll discuss how devices and networks can—and should be—protected as part of an overall strategy of sensor deployment, data collection and analysis.

The Industrial Control Systems (ICSs), which preceded today’s IoT systems, existed in silos and used proprietary protocols, networks and technology. Conventional network and system security were provided by establishing a perimeter around the entities which needed to be safeguarded. Once established, firewalls, intrusion detection and prevention systems served as the foundation for security and virtual private networks provided a tunnel into the network. But IoT is drastically reshaping how applications and ICSs operate and are secured.

Current IoT devices have ubiquitous connectivity to the network using open standards, which no doubt is beneficial in providing rich functionality, but enhances threat surface significantly. The network, the things, the associated systems and applications need to be secured. IoT induced reshaping is based on the differentiators in an IoT network versus conventional ones. The differentiators individually impact security, but, when combined, portend exponential security threats and subsequent impacts. Key IoT differences include:

- (1) The network topology is ever expanding, pushing the boundaries of the network with the constant addition of functionalities, applications, devices and equipment.
- (2) With the constant network expansion, more functionality gets incorporated in to the edge of the network. While edge computing has many advantages, including improved response time and localized services and processing, it exposes the network as a whole to significant vulnerabilities.
- (3) Diverse functionalities/applications require different levels of authorization and access to data and systems in the network, edge and backend systems. Improper use of authorization and access or providing blanket permissions to a large number of systems may make the entire network vulnerable.



- (4) There is shared multi-tenant cloud usage for compute and storage. Vulnerability of any service or user in any part of the network can impact the entire network.
- (5) Critical infrastructure in factories, utilities, cities, etc., can be managed and operated from mobile consoles and personal devices, which brings additional security challenges.
- (6) Due to perpetual addition of new applications and functionalities, the network is in a continuous state of flux.
- (7) State, federal and international data privacy and security regulatory compliance requirements, under the Federal Trade Commission's (FTC) privacy framework and Federal Trade Commission Act ("FTC Act"), the California Consumer Privacy Act (CCPA), the European Union's (EU) General Data Protection Regulation (GDPR), city/state specific regulations etc., mandate operators/utilities to ensure data security and privacy over a broad range of personal information and/or data.

Ultimately, the highest level of security is paramount, and the network is only as secure as the weakest connected device, with the quantity of devices expanding and changing daily (theoretically). Furthermore, conventionally used perimeter-based security measures (e.g., firewalls, De-Militarized Zones (DMZs), etc.) are not sufficient for IoT networks as one has to consider both fixed IoT devices (e.g., water meters) and mobile IoT devices (fleet management) that extend the edge and make it dynamic. For example, government employees leveraging mobile IoT devices frequently access critical infrastructure networks in cities and utilities. Malicious access through one of these IoT devices could cause catastrophic network effects.

Compounding security problem are data flows (from sensors through networks to public clouds and third-party devices and services) outside of service provider or network operator control. This adds additional dimensions for IoT data security. For instance, data exchanged with water meters by a utility over an operator's network needs to be secured and protected for privacy across all the segments during transit and storage. It requires clear delineation of responsibilities and adherence to protocols for secure operation and management of the devices in transit and data generated from these devices. Not only is the IoT infrastructure itself exposed, but citizens' personal information (PI) could also be exposed, which could be out of compliance with data security and privacy regulations such as California Consumer Privacy Act (CCPA) and General Data Protection and Regulation (GDPR).

According to some industry studies, IoT security still comes up as the number one deterrent to IoT adoption year after year.<sup>[1]</sup> Having mentioned different vulnerabilities of an IoT network and its impact on network security, we'll now discuss the approaches that may be adapted to mitigate some of these concerns.

IoT devices are relatively inexpensive, have ubiquitous connectivity to the critical network infrastructure, possess enough compute and storage, but probably are not ruggedized enough from a security perspective. This makes them attractive for rogue players with malicious intent to potentially harm the network systems, applications and critical infrastructure. A case in point—in 2016, in one of the infamous IoT DDoS attacks, a botnet infected nearly 65,000 devices in its first 20 hours, doubling in size every 76 minutes.<sup>[2]</sup> Before eradicated, it infected thousands of devices globally and halted the Internet for a portion of the U.S. for some time.

Cloud computing has become pervasive with many domains, including IoT, rapidly adapting to it. Cloud computing is attractive in providing flexibility in deployment, providing scalability and is cost effective. Also, the multitenancy aspects of cloud are especially attractive as it is an enabler for scaling different applications or verticals of the IoT network independently. However, as some recent incidents<sup>[3]</sup> across many organizations, including the US Navy, major corporations and more than a dozen cloud providers globally have shown, vulnerability in one segment can be exploited to impact other parts of the network. Even if one maintains one's own network diligently from a security perspective, it may still become a victim of vulnerability of a different entity in the shared cloud. The sensitivity level of the data may trigger regulatory compliance requirements or create enhanced risk if not properly secured under the FTC's privacy framework or GDPR. Private and hybrid clouds are also not immune from this problem by virtue of their accessibility needs over a public network.

Having mentioned different vulnerabilities of an IoT network and the impact on network security, now let us discuss the approaches that may be adapted to mitigate some of these concerns.

## 2. Proposed IoT Security Framework

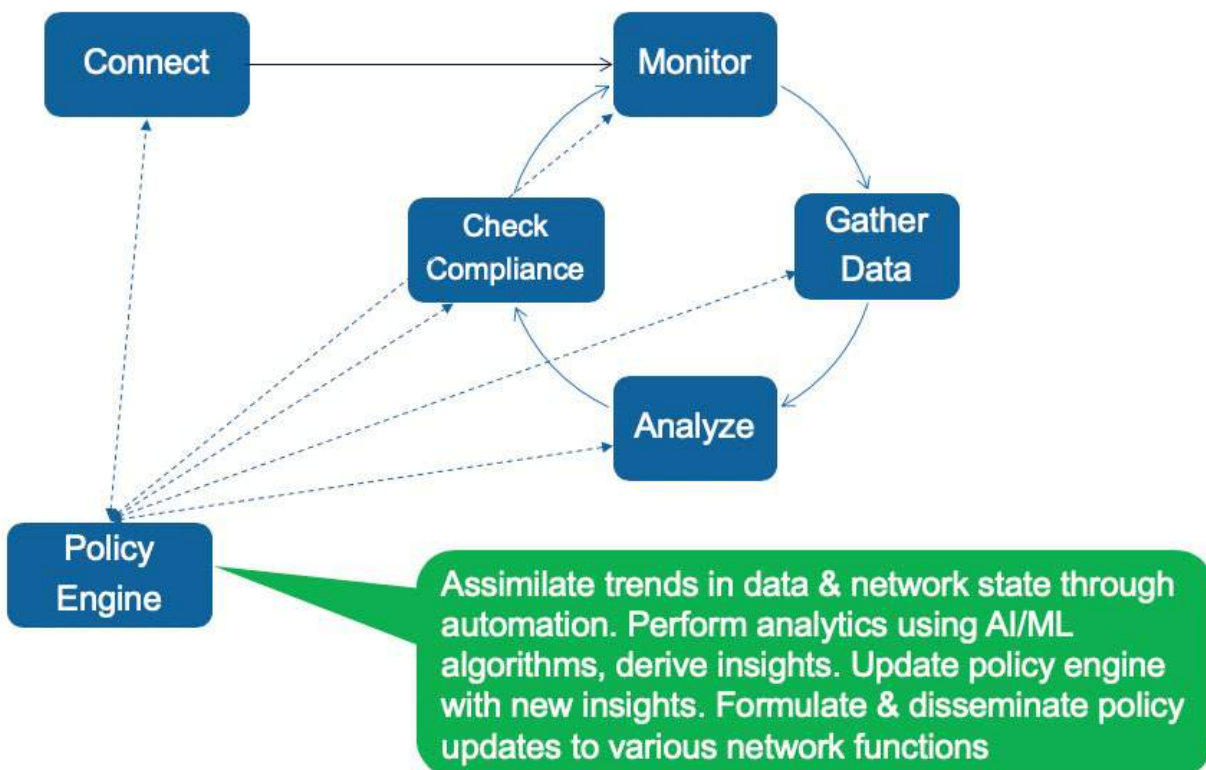


Figure 1 - IoT Security Life Cycle

As depicted above, each entity in the IoT network should be provided with *minimum viable access* (to connectivity, bandwidth, amount of data transmitted and frequency of transmissions, authorization to connect to systems and access to data). All this should be continuously monitored through an automated process to gather the data, analyze to discern insights and identify anomalies. The policies need to be updated based on gathered insights, observed trends in data and operator input. This minimal viable access and compliance with policies for each entity, including devices, edge compute nodes or more complex core/access network functions, forms the foundation of IoT network security.

To illustrate the importance of this point, it may be noted that in 2017 a casino lost its high roller database through the network connectivity provided for a fish tank.<sup>[4]</sup> The fish tank had sensors connected to a PC that regulated the temperature, food delivery and cleanliness of the tank. This vector was used to steal the high roller database. Needless to say, a temperature control system for a fish tank does not need, and should not have had, any possible connectivity to critical data, and it is a failure of formulating and enforcing effective security policies (e.g. minimal viable access to connectivity, data and authorization).

It is important to recognize that, network security is **not** an isolated standalone function, rather it is an overarching, all-encompassing characteristic of a system. As the saying goes, a given system is only as secure as its weakest constituent component. It is imperative that the security of the entire system be looked into as a whole, not as individual isolated components.

## **2.1. Connect:**

The diversity and volume of IoT use cases are numerous. Equally large are associated devices, network components and diversity of their connectivity needs. This diversity certainly increases the threat surface area, variety of threat vectors and vulnerabilities. To mitigate this risk, a strong connectivity policy needs to be created and enforced. The policy should be customizable and consider the characteristics of each device, the evolving trends in the network, the availability of the network resources, the relative priority and criticality of various functions and their connectivity needs. This can be arrived at after a thorough analysis of various components in the network and building a subsequent enforcement framework.

## **2.2. Monitor:**

The diversity of devices and network functions in an IoT network have different capabilities and provide different metrics to gauge and monitor these capabilities. The monitoring function needs to consider a profile for each entity (devices, edge/core/access network functions) and create a set of characteristics to be monitored and adapt these characteristics to evolving conditions in the network. Some of these may be conditional on meeting certain thresholds in different areas.

Continuous monitoring of the different aspects of the network, such as traffic patterns, directions of data flow, any norm breaking trends in data or traffic, and evolution in the network are key to highlighting the existence of potential security threats and identifying them. Today's network technologies also offer greater visibility into application and device activity. Software designed to detect anomalous behavior at the network level, revealing Distributed Denial of Service

(DDoS) and other attacks, can now leverage artificial intelligence (AI) and machine learning (ML) to respond. Changes in behavior at the application and device level can raise alerts.

### **2.3. Analyze:**

The IoT network is diverse in topology and in its constituent network elements. The security issues often cannot be detected or identified by looking at the snapshot of the network at any given time in isolation. It requires a thorough analysis of data gathered from the above mentioned continuous monitoring of the different aspects of the network. The analysis needs to examine not only individual snapshots of the network, but also needs to correlate data points from different parts of the network and across different points in time to identify any emerging trends and discern insights. The insights derived may lead to the addition of new policies or updating existing policies.

### **2.4. Compliance:**

This is a gating function, which ensures all entities in the network adhere to respective policies and rules that need to be enforced strictly without exception. Any updates needed to these enforcement rules should come through the policy engine after careful analysis of the impact of the proposed changes across the entire network based on available data and/or operator input. Automation of policy adoption and enforcement is critical for ensuring continued compliance and building resiliency into the process. Automation also helps recognize and address shifting trends in the network.

### **2.5. Improvement cycle:**

The IoT Security Life Cycle is a continuous improvement cycle. The accumulation of data and analytics previously mentioned and the correlation of different data points across time and different parts of network provide insights about the network and applications in their current state, as well as for emerging trends in the future. These insights can be used to mitigate current threats and plan for, and address, any emerging threats in the future even before these are materialized.

The improvement cycle comprises identifying potential threats and fine-tuning security policies to adapt to the perceived and emerging threat vectors, based on the insights gathered from analysis and correlation of various data points. This process of keeping the security policies in sync with current and emerging needs of the network allows for continuously updating the policy engine based on the insights. This can also make the network more efficient by eliminating any redundancies, in addition to enhancing the robustness of the network. In summary, this continuous improvement process helps to mitigate security threats of the IoT network, constituent devices, supported applications and enables realization of value of five Vs (*volume, velocity, variety, veracity and value*) of IoT data generated and processed.

## **3. IoT Data : Security, Integrity and Ownership**

People now often say that data is the new oil. However, it doesn't fully characterize the value of IoT data. To quote Adam Schlosser of the World Economic Forum, "unlike oil, the value of data

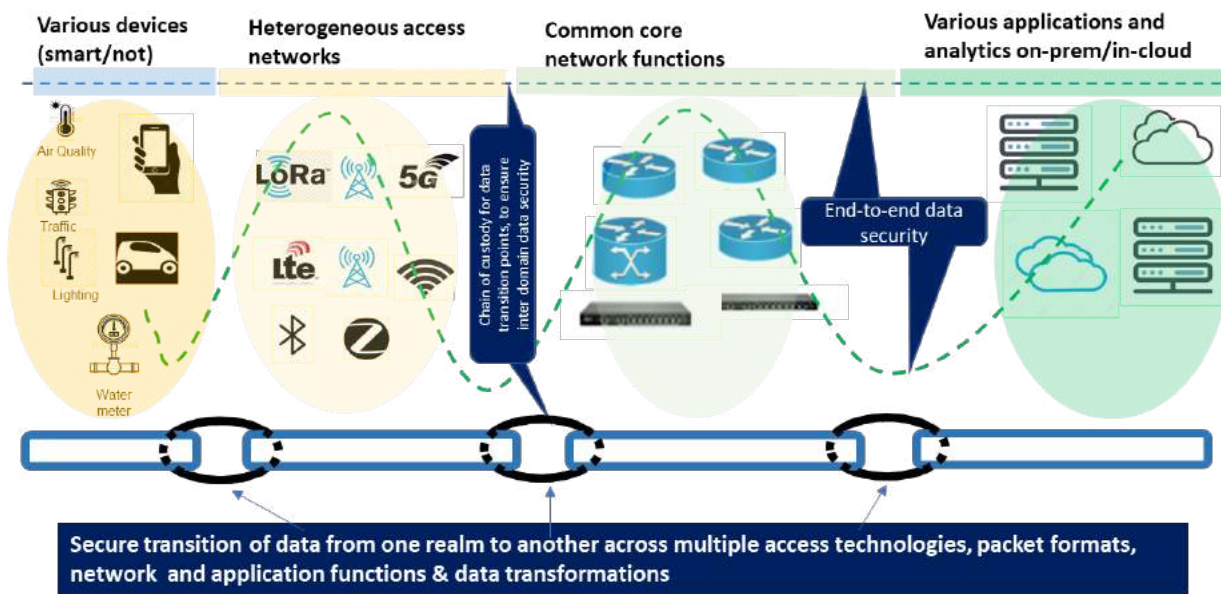
doesn't grow by merely accumulating more. It is the insights generated through analytics and combinations of different data sets that generate the real value.”<sup>[5]</sup>

IoT data promises to convey more useful information than ever before, and the volume and velocity could improve the speed and accuracy of all sorts of business and strategic decisions significantly. Reliability and security of data is essential to make the available information actionable. IoT data security, integrity and ownership are the building blocks of this trust and are of paramount importance.

### **3.1. Security and Integrity:**

As described previously, IoT networks comprise a diverse collection of devices, and applications on the network are evolving and generating enormous amount of data. The security solution for this dynamic, ever-evolving network cannot be static. The security policies, framework and measures need to adapt to this diversity in devices and applications. We propose the following framework for securing IoT data and ensuring its integrity:

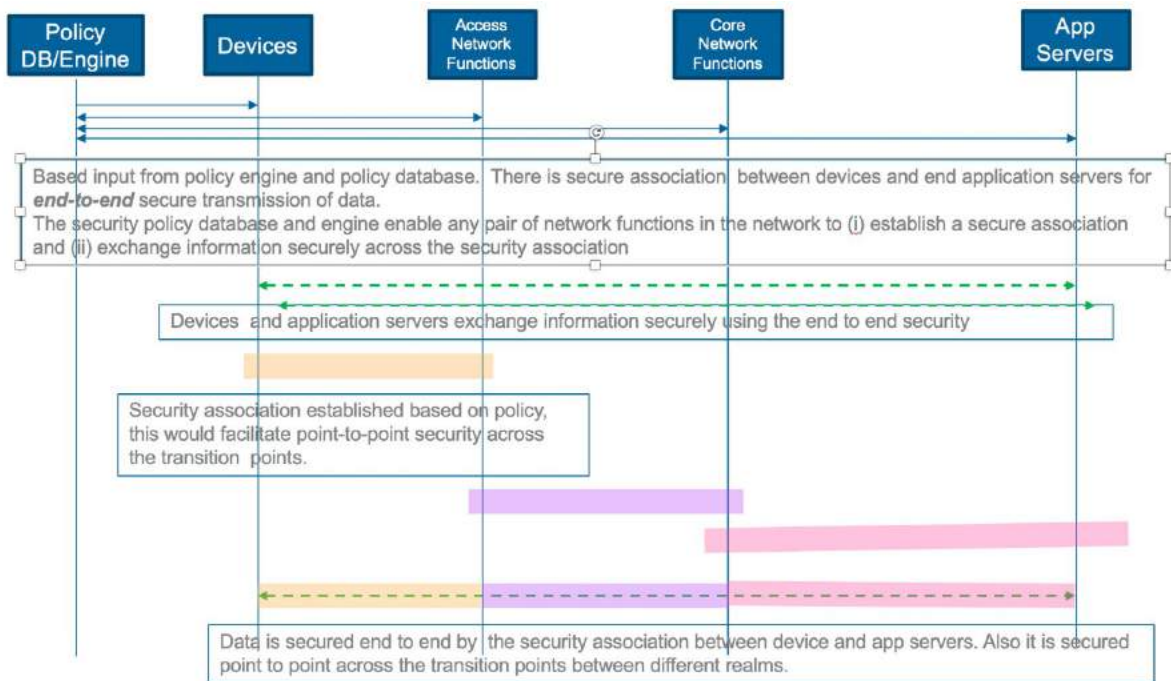
- The data shall be secured while it traverses through the network across different network elements, including end to end.
- A chain of custody shall be established for the data generated from the point of origin until it is processed/transmitted/stored.
- There may be different stakeholders for each segment of the network, such as device-users, access network providers, core network providers, application server providers, etc.
- There are clearly delineated responsibilities and expectations for each of the stakeholders regarding how to handle the data as it enters the entity and how the data gets processed and leaves the entity.
- Any two entities exchanging data shall have a *security association (SA)*, *minimum viable connectivity*, *access and authorization driven by the security policy*, which may differ in each direction of data transfer.
- The overhead associated with above mentioned security association and chain of custody is not cumulative and additive for each piece of data transferred across this interface.
- All data is secured in transit and storage. The data security shall include both ciphering and integrity protection as dictated by the security policy.
- Each entity in the network shall have no more visibility and authorization than absolutely necessary to perform its function. This applies to data as well as other network resources.
- AI/ML based algorithms shall be used to correlate different metrics in the network to detect and mitigate any suspicious activity.
- A process is established to isolate and quarantine the impacted applications, devices and segments of the network through a rapid response system.
- A security model shall be adopted to create micro-segments and enable granular enforcement of security policies. This model shall be used to manage data and devices, and migration of data across different components, both external and internal



**Figure 2 - Chain of Custody for IoT Data Security**

To establish this chain of custody (CoC) for the data as it traverses through the network across different network elements, access/segment specific methods shall be used. For resource constrained (i.e. low bandwidth, smaller payloads, high latency) segments of the network, access segment specific methods are needed. Any two entities exchanging data shall have a *secure association, minimum viable connectivity, access and authorization driven by the security policy*. For the segments of the network, based on the sensitivity of the data and availability of bandwidth, a blockchain-based method may be adapted. It may be noted that the distributed ledger/blockchain based methods are suitable for IP networks and have cumulative overhead for securing the data. Our proposed method is adaptable to different access technologies and transport methods and does not incur cumulative overheads. However, the proposed chain of custody is complementary to blockchain-based technologies. To secure data in any given segment of the network, blockchain-based methods may be adapted within this framework. A robust set of security policies shall be formulated taking the above into account. The more granular these policies are, the more fine-grained control it provides on different aspects of data, applications and network. Automation is key to ensuring persistent compliance with established security policies with aforementioned characteristics.

Below is an illustration of messages and information exchanged between different entities in the network to establish a policy driven chain of custody for secure exchange of data without incurring cumulative overhead.



**Figure 3 - Message Flow for Chain of Custody for IoT Data Security**

### 3.2. Ownership:

Data is an asset. With IoT, data is collected and acted upon at different points within the larger IoT framework. The stakeholders for different parts of the framework may be different. For this reason, it is essential that ownership of different aspects of the data and its visibility among the stakeholders is clearly defined upfront.

As an illustration, consider autonomous/connected vehicle data. Vehicle manufacturers, the companies that produce the individual components, telecommunications providers and possibly insurance companies may all want the data produced during an operation. Meanwhile, the vehicle owner might have qualms about sharing personal data. In any event, there's value in that data, and, naturally, each stakeholder desires to capitalize on it.

With Smart Cities, data ownership is further complicated. In addition to different stakeholders like utilities, consumers, etc., city projects may have regulatory compliance requirements, as well as implications associated with public funding of city projects. Unlike the enterprise data, public funding of various city projects often triggers various concerns of the public's right to information aspects for the city's IoT data. Usage, disclosure and monetization of this data may have additional constraints. To avoid any confusion about the ownership of data and insights derived from it, it is imperative to establish upfront roles, responsibilities and rights as to the ownership of data among all stakeholders at the beginning of the project.



## 4. Standards support

Technologies like IoT evolve faster than related standards. What makes IoT unique is that standards are required for multiple enablers such as communications, semiconductors, devices, privacy and security to name a few. Standards are also the foundation for interoperability. The faster they can be defined and adopted, the faster IoT systems will deliver potential advantages. The IoT standards for access networks, security and core functions are at various stages of evolution.

In a report titled *Hype Cycle for IoT Standards and Protocols, 2020*, <sup>[6]</sup> Gartner states “we still see many standards overlapping, or competing directly, especially in areas with large commercial potential. There are also areas of the IoT in which standards are incomplete or lack full stack support. Consequently, new standards will continue to emerge in the coming years.” Currently available security standards from various standards bodies such as LoRa Alliance, 3GPP, IETF, NIST, GSM-A, etc., can be leveraged, adapted and enhanced to secure an IoT network. These standards, originating from diverse standards bodies, were originally intended to address different needs. Adaptation of these for securing an IoT network requires a thorough analysis of the network being secured, configuring, customization and integrating these protocols together to address the specific needs of the IoT network under consideration.

## 5. Conclusion

There are many challenges in the rapidly evolving IoT applications, network and devices with huge amounts of data generated. Security is of paramount importance among all of these. To realize the promise and potential of the IoT network, it is clear that the data generated and needed to be acted upon must be secured and trusted. Ownership is an imperative and requires clear and upfront delineation of roles and responsibilities of all stakeholders.

## 6. Acknowledgements

The Authors would like to express sincere gratitude to Mr. Satya Parimi and Ms. Patricia Zullo for their time and effort for many cycles of reviews, industry insights and valuable suggestions for improvement.



## Abbreviations

|       |                                                |
|-------|------------------------------------------------|
| 3GPP  | 3rd Generation Partnership Project             |
| AI    | Artificial Intelligence                        |
| CCPA  | California Consumer Privacy Act                |
| CoC   | Chain of Custody                               |
| DoS   | Denial of Service                              |
| DDoS  | Distributed Denial of Service                  |
| GDPR  | General Data Protection and Regulation         |
| GSM-A | Groupe Speciale Mobile(GSM) Association        |
| IETF  | Internet Engineering Task Force                |
| IoT   | Internet of Things                             |
| LoRa  | Long Range                                     |
| ML    | Machine Learning                               |
| NIST  | National Institute of Standards and Technology |
| PI    | Personal Information                           |
| SA    | Security Association                           |

## Bibliography & References

1. “IoT Security is Still a Major Barrier to Adoption,” Kathryn Weldon, Global Data, <https://itcblogs.currentanalysis.com/>, March 21, 2019
2. “Inside the infamous Mirai IoT Botnet: A Retrospective Analysis.” Elie Bursztein. Cloudflare.com, December 14, 2017. <https://blog.cloudflare.com/inside-mirai-the-infamous-iot-botnet-a-retrospective-analysis/>
3. “Ghosts in the Clouds: Inside China’s Major Corporate Hack”, Rob Barry and Dustin Volz. Wall Street Journal, December 30, 2019: <https://www.wsj.com/articles/ghosts-in-the-clouds-inside-chinas-major-corporate-hack-11577729061>
4. “How a fish tank helped hack a casino,” Alex Schiffer, Washington Post, July 21, 2017: <https://www.washingtonpost.com/news/innovations/wp/2017/07/21/how-a-fish-tank-helped-hack-a-casino/>
5. “You may have heard data is the new oil. It's not,” Adam Schlosser, World Economic Forum, Jan 10, 2018, [www.weforum.org/agenda/2018/01/data-is-not-the-new-oil/](http://www.weforum.org/agenda/2018/01/data-is-not-the-new-oil/)
6. Gartner, Inc. “Hype Cycle for IoT Standards and Protocols, 2020,” Bill Ray, June 30, 2020

# **Decoding The Bandwidth Surge During Covid-19 Pandemic**

## **An In-depth Study On DOCSIS Upstream Bandwidth Surge And Their Impact On Video Conferencing**

A Technical Paper prepared for SCTE•ISBE by

**Ram Ranganathan**  
Director, Office of CTO  
COMMSCOPE  
Mississauga, CANADA  
905 568 7317  
ram.ranganathan@commscope.com

**Chris Markovich**, COMMSCOPE

**Tushar Mathur**, COMMSCOPE

**Omar Abu-Hijleh**, COMMSCOPE

**Thomas Cloonan**, COMMSCOPE

**John Ulm**, COMMSCOPE

# Table of Contents

| <b>Title</b>                                                    | <b>Page Number</b> |
|-----------------------------------------------------------------|--------------------|
| 1. Introduction.....                                            | 3                  |
| 2. Impact of bandwidth surge on DOCSIS networks .....           | 3                  |
| 3. The Experiment.....                                          | 4                  |
| 3.1. Focus on Video Conferencing Applications .....             | 5                  |
| 3.2. Experiment Goals .....                                     | 5                  |
| 3.3. Experiment Approach .....                                  | 6                  |
| 3.4. Inputs .....                                               | 7                  |
| 3.5. Experiment Setup .....                                     | 8                  |
| 4. Observations and Analysis.....                               | 9                  |
| 4.1. Network Congestion related measurements.....               | 9                  |
| 4.2. Understanding Video Conferencing App Operating Points..... | 10                 |
| 4.3. Speed test & Conferencing App throughput behavior .....    | 11                 |
| 4.4. Video Conferencing Latency measurements .....              | 13                 |
| 5. Conclusion .....                                             | 16                 |
| Abbreviations.....                                              | 16                 |
| Bibliography & References .....                                 | 17                 |

## List of Figures

| <b>Title</b>                                                                                          | <b>Page Number</b> |
|-------------------------------------------------------------------------------------------------------|--------------------|
| Figure 1 – Downstream and Upstream impact during COVID-19.....                                        | 4                  |
| Figure 2 – Testbed Setup .....                                                                        | 7                  |
| Figure 3 – Setup 1: QoE impact due to DOCSIS upstream.....                                            | 8                  |
| Figure 4 – Setup 2: QoE impact due to in-home and channel congestion .....                            | 9                  |
| Figure 5 – Latency, Jitter and Packet Loss under DOCSIS upstream load.....                            | 9                  |
| Figure 6 – App 1 without (left) and with (right) heavy network load .....                             | 10                 |
| Figure 7 – App 2 without (left) and with (right) heavy network load .....                             | 11                 |
| Figure 8 – App 3 without (left) and with (right) heavy network load .....                             | 11                 |
| Figure 9 – Speed test along with multiple conference apps .....                                       | 12                 |
| Figure 10 – Speed achieved by a single conference session under CM+channel load.....                  | 13                 |
| Figure 11 – Mean Latency for conferencing apps under channel loading ONLY .....                       | 14                 |
| Figure 12 – Mean Latency for conferencing apps under CM+channel load .....                            | 14                 |
| Figure 13 – 95 <sup>th</sup> percentile latency for conferencing apps under channel loading only..... | 15                 |
| Figure 14 – 95 <sup>th</sup> percentile latency for conferencing apps under CM+channel load .....     | 15                 |

## List of Tables

| <b>Title</b>                                             | <b>Page Number</b> |
|----------------------------------------------------------|--------------------|
| Table 1 – Experiment Parameters.....                     | 8                  |
| Table 2 – Conferencing App operating “sweet spots” ..... | 10                 |

# 1. Introduction

During the COVID-19 pandemic, millions of people around the globe have become far more reliant on their broadband internet connection to stay connected to their family, friends, and co-workers. The broadband internet has emerged as an essential technology that is keeping society together and the economy running in this stressful time.

The growth of broadband internet traffic was bound to happen due to “Work from home” and “Stay at home” instructions around the world during the pandemic. The dependence upon virtual meeting and online collaboration applications has become increasingly important during the pandemic as well. The web apps that are helping us stay connected must work reliably for people to be able to socially connect and stay productive. A few key questions had to be answered quickly during these crises by the broadband operators and vendors – Are the bandwidth surges during the COVID-19 pandemic breaking the broadband Internet connectivity? How are the commonly used applications handling these surges?

There are many popular web applications that drive the conferencing audio and video traffic in any broadband network. Many cable operators have reported a sudden growth of Upstream and Downstream traffic in their networks resulting from these conferencing apps<sup>1</sup>. Due to the tighter constraints that currently exist on the Upstream spectrum, the DOCSIS Upstream is more sensitive to this uptick in bandwidth consumption. In this paper, we will present our research and analysis of the major sources of DOCSIS Upstream bandwidth surge. Specifically, this paper will provide insight into the surge impacts on video conferencing applications and observations on latency, jitter, packet loss, and throughput as DOCSIS Upstream channel utilizations reach an extraordinary tipping point.


## 2. Impact of bandwidth surge on DOCSIS networks

Key application drivers for the bandwidth surge are quite similar on both DOCSIS upstream and downstream networks, including:

- Work-at-home applications – VPN traffic, video conferencing (Cisco Webex™, Microsoft Teams™, Zoom™ and others) and cloud sync for upstream traffic
- Remote learning applications including video conferencing/Zoom™
- Video streaming (Netflix™, Hulu™, Amazon Prime™, Disney+™ and others) predominantly as downstream traffic
- Gaming applications (more actual gaming traffic on the upstream as well as Twitch broadcasts, while downstream is impacted by Twitch video streaming of games)
- Social networks (Facetime/Skype increased usage between families etc.)
- Increased TCP Acks on the upstream due to increased downstream traffic

Based on our studies, as shown in Figure 1, downstream bandwidth usage increased by 12-25% during the busy hours (typically evening prime-time) and also significantly increased during midday by 26-86% again due to video conferencing related traffic loads. On the upstream, this surge was even more profound as peak periods of upstream bandwidth shifted towards the afternoons and the overall “sleep or quiet” periods of traffic was much lower. Upstream peak usage during typical evening peak periods increased by 20-50%, while midday usage showed a much more substantial increase of 30-150%.

The increased upstream usage is typically an achilles heel of DOCSIS deployments for Multiple System Operators (MSOs) as today’s deployment are based on low-split with upstream channels in the frequency range of 20-42MHz or 20-65MHz. Future migration to mid or high splits may reduce the impact of such an upstream bandwidth spike.



**Figure 1 – Downstream and Upstream impact during COVID-19**

Based on data analytics collected from multiple MSOs using the CommScope ServAssure program, we were able to observe network utilizations levels before and after the pandemic<sup>2</sup>. DOCSIS channel utilization increased significantly week over week during the pandemic lock-down for both upstream and downstream. While these rapid surges may have pulled back as the lockdown eased, the overall congestion observed, and the significant number of service groups impacted shows the need to be prepared for such bandwidth surges in the future. This paper focuses on impact of upstream DOCSIS channel congestion on applications such as video conferencing. It also studies how key performance metrics (KPIs) such as latency and throughput perform under such constrained channel conditions.

Overall, DOCSIS based cable broadband networks managed to handle the bandwidth surge pretty well. Key factors contributing to this positive result include:

- CMTS scheduling algorithms are adaptive handling congestion with fairness based algorithms
- Most Internet applications, especially high bandwidth ones like streaming video and conferencing apps, are adaptive in nature and adjust well to bandwidth fluctuations
- Mechanisms in modems and gateways like Active Queue Management have worked in tandem with application adaptive behavior to keep latency and jitter to somewhat an acceptable level
- Several MSOs have planned additional capacity headroom in their network planning and deployment (for example, 1.2 times Tmax rate to absorb SLA surges)

However, it should be noted that it was not all positive news when it comes to the sudden increase in bandwidth need. Applications like gaming suffered from increased “lag” (latency/jitter) due to buffer bloat issues. Due to packet delays and drops, video tiling and macroblocking effects were also observed in streaming and video conferencing applications. In addition, for subscribers with data caps for their broadband access, those data caps were hit more frequently during the pandemic.

### 3. The Experiment

Our study in this paper focuses specifically on video conferencing applications and how they are impacted by increased in-home traffic and by shared traffic within a DOCSIS service group. In a typical cable broadband access network, congestions points are typically encountered in these areas:

- In-home Wi-Fi Congestion: Predominantly caused due to large number of WIFI users or endpoints in the home. Other factors like wireless interference and distance-to-the-WIFI-gateway related issues also has an impact to overall network congestion, latency and packet losses.

- Cable Modem (CM) Congestion: This is often due to increased bandwidth from in-home applications. Algorithms like Active Queue Management on CMs that manage buffer bloat also influence performance KPIs.
- DOCSIS Channel Congestion: This is attributed to increased bandwidth usage in the neighborhood (associated with the specific Service Group) and scheduling algorithms in the Cable Modem Termination Systems (CMTS) that go hand-in-hand in managing this increased bandwidth.

The focal point of the study presented here is on #2 and #3 from the list above, specifically on the DOCSIS upstream channel which is usually more constrained due to limited bandwidth availability.

### **3.1. Focus on Video Conferencing Applications**

As noted earlier, one of the key drivers of Internet traffic during the COVID-19 bandwidth surge was video conferencing. Applications such as Cisco Webex™, Zoom™ conferencing, Google Meet™ and Microsoft Teams™ were used extensively by work-at-home employees, remote learning and other e-learning activities and even as virtual video meetups for social events.

For the sake of this study, we took three of such commonly used applications which allow simultaneous sharing of audio and video along with the ability to be in “presentation mode” with a number of users. The chosen applications will be referred to as App 1, App 2 and App 3 in the remainder of this paper as the intention of our study is not critique the behavior of such applications, but rather to illustrate how diverse the behavior can be and how changing bandwidth conditions can affect the QoE of various conferencing apps.

All three video conferencing apps chosen are capable of running over UDP or TCP as the transport protocol for media traffic, however the default protocol (and the preferred one) is typically UDP. For our experiments, all conferencing apps were using UDP as their transport protocol for media (audio/video) traffic.

It should be noted that outside of audio and traffic, each conferencing application also had its own set of signaling protocols. Some of these leverage existing protocol standards such as Real-time Transport Control Protocol (RTCP) and Session Traversal Utilities for NAT (STUN). For all experiments, we did not make any modifications to the behavior of these signaling protocols but some of the traffic patterns resulting from how these applications reacted to packet loss and latency via signaling were observed.

It was also determined that all three apps have their own flavor of “adaptive streaming” behavior to handle congestion control and ensure that the end user behavior is kept to a minimum due to changing network conditions (and other factors in the streaming path).

### **3.2. Experiment Goals**

Some key questions we attempt to find solutions for include:

- 1) How do key performance metrics like latency, jitter and application throughput change with increased DOCSIS channel utilization and in-home congestion?
- 2) What happens to Quality of Experience (QoE) of video conferencing applications when DOCSIS upstream channel is congested (70% or higher capacity used)?
- 3) How does the QoE change when the congested channel issue is combined with increasing home Internet traffic including multiple conferencing sessions?

### 3.3. Experiment Approach

Our initial setup in conducting these experiments was to determine how many active cable modems or gateways are “active and interacting” in a typical DOCSIS Service Group during the course of the pandemic. We expected that our previous figures on such active CMs may not be accurate due to the increased activity due to the lockdown effect in the pandemic. Our study of traffic patterns from North American MSOs during the late March and early April timeline showed us that typically 75-80% CMs were active and transmitting during a 15-second window during peak traffic periods. This number (80%) was used in our study to emulate traffic from multiple modems in a service group.

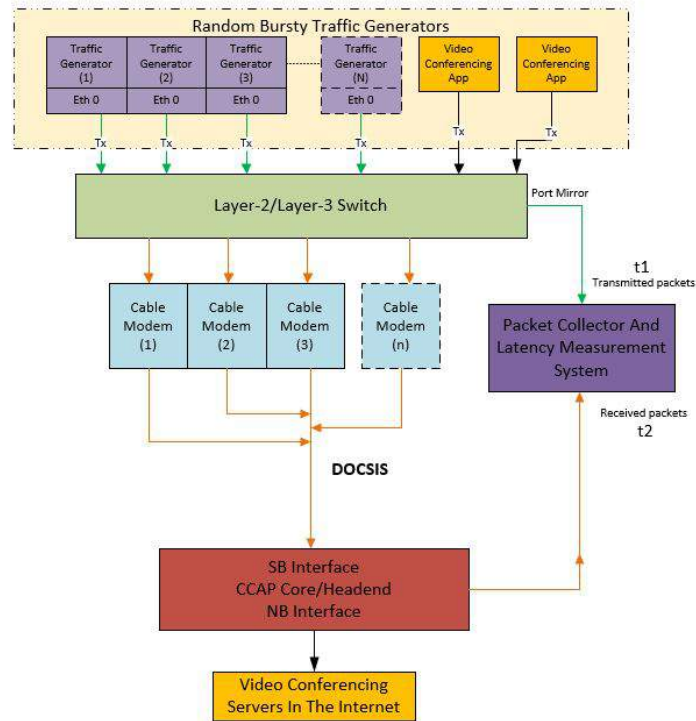
Another key aspect that was considered during these experiments was using emulation versus simulation techniques. Due to the complex nature of traffic patterns and changing server behavior of conferencing applications, it was decided that real-world traffic emulation communicating with real-world conferencing servers was the preferred approach.

As part of the experiment setup, we had an array of traffic generators that was modelled from statistical models derived from real-world, sampled DOCSIS upstream traffic. Specifically, bursty traffic loads were initiated from these traffic generators using statistical models of three key parameters:

- Duration of traffic burst
- Varying bandwidth levels within each traffic burst
- Inter-burst interval – idle time between successive bursts

The DOCSIS upstream was congested to the desired channel congestion levels (70%, 80% etc.) depending on the experiment needs using these traffic generators. Once this was achieved, real-world video conferencing applications were run with multiple users with the target user on a PC behind the CM under test. The test PC was then broadcasting a video stream and the QoE of that media stream was observed. Additional metrics such as latency and throughput were measured during the video conferencing session.

A conceptual representation of the experiment testbed is shown in Figure 2.



**Figure 2 – Testbed Setup**

### 3.4. Inputs

A service group size of 400 cable modems was chosen for the experiment setup. Based on the prior analysis of real-world subscriber activity, 80% of those modems were considered to be active and interacting.

Four service tiers were part of the testbed as shown in Table 1, along with the distribution of CMs across these tiers. 319 of the 320 modems (excluding the CM under test) were then connected to Traffic generators (as Customer Premises Equipment/CPE devices) and real-world modelled bursty traffic were sent from these generators based on the experiment's need to achieve the desired DOCSIS upstream channel utilization (70%, 80% and so on).

The “**CM under test**” was chosen from service Tier 3 which had a 25Mbps downstream and 5Mbps upstream Service Level Agreement (SLA). This was chosen as the “minimal acceptable upstream bandwidth” needed to run at least a couple of video conferencing applications from the home. Depending on the experiment, devices behind the test CM were one or more PCs running video conferencing applications or traffic generators for throughput/latency/packet loss measurements.

For all video conferencing tests, the test video was chosen to be an animated GIF image (repetitive motion) to account for motion based video transmission and the setup was maintained in a controlled lighting environment with minimal human interference during tests. Fluctuations in video quality (lighting, scene changes etc.) can result in significant variation in transmission bandwidth and the purpose of the controlled environment was to keep any drastic bandwidth fluctuations to a minimum. It should be noted that the animated GIF image was streamed via video during the conference rather than as a “screen



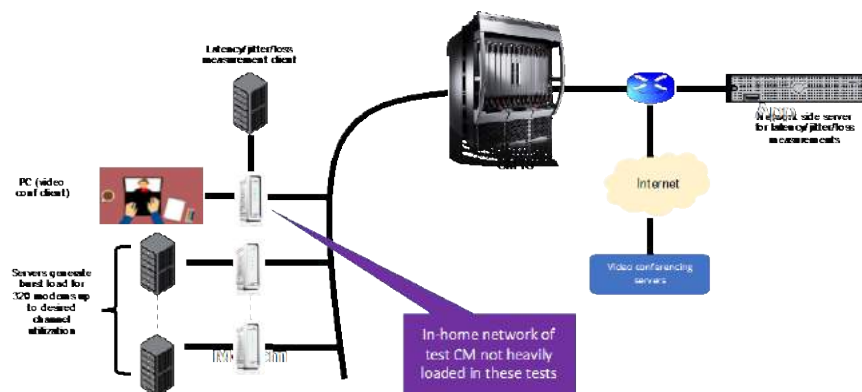
share”. Screen sharing typically uses less bandwidth than video transmission itself in such applications, so the experiment was focused on the more bandwidth-intensive video stream sharing.

**Table 1 – Experiment Parameters**

| Key parameters                   | Values                                                                                                                    |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| Number of active CMs in SG       | 320                                                                                                                       |
| Service tiers (DS/US bandwidth)  | Tier 1: 100Mbps DS/40Mbps US<br>Tier 2: 50Mbps DS/20Mbps US<br>Tier 3: 25Mbps DS/5Mbps US<br>Tier 4: 12 Mbps DS/1 Mbps US |
| Distribution of service tiers    | Tier 1: 10%<br>Tier 2: 65%<br>Tier 3: 15%<br>Tier 4: 10%                                                                  |
| CM configuration details         | DOCSIS 3.1 CMs, default values for AQM/buffer controls, ACK suppression enabled                                           |
| DOCSIS Upstream Channel capacity | ~100Mbps (4xATDMA channels: 3x6.4Mhz channels and 1x3.2Mhz channel)<br>DS is not congested in these tests                 |

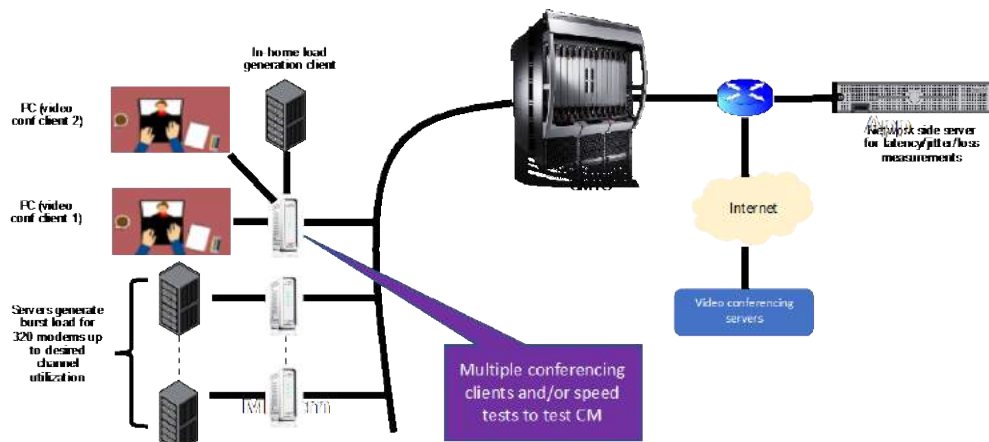
### 3.5. Experiment Setup

Two primary setups were used for these experiments. The first set of experiments were focused on DOCSIS upstream channel loading using traffic generators as shown in Figure 3. Along with the traffic generators loading the DOCSIS channel, a single PC client is behind the test CM on a video conferencing session with 4 other participants. A small “test” stream is used to measure latency and packet loss metrics due to the congestion.



**Figure 3 – Setup 1: QoE impact due to DOCSIS upstream**

The second setup, as shown in Figure 4, adds multiple conferencing clients behind the test CM and also generates in-home traffic loads to emulate aspects like a speed test or file upload. Overall throughput achieved by the conferencing clients as well as other KPIs were measured using this setup.




**Figure 4 – Setup 2: QoE impact due to in-home and channel congestion**

## 4. Observations and Analysis

### 4.1. Network Congestion related measurements

To obtain an understanding and a good baseline of how key performance metrics like latency, jitter and packet loss behaved under network load, experiments were conducted loading the DOCSIS upstream with bursty “real-world-like” emulated traffic from the traffic generators. Test streams were injected behind cable modems in every service tier defined in Table 1 to measure (mean) latency, jitter and packet loss at three key channel loading points – 78%, 100% and 122%. It should be noted that this baseline experiment did not have other high bandwidth applications like video conferencing.



**Figure 5 – Latency, Jitter and Packet Loss under DOCSIS upstream load**

From Figure 5, it can be observed that (mean) upstream latency was not impacted heavily by channel utilization unless the Tmax was at the lowest performing tier (=1Mbps). The jitter graph shows that even with 5Mbps upstream Tmax, packet jitter starts becoming an issue as channel gets congested. This was further validated during our video conferencing tests which we describe in detail in the sections below. Packet loss in general was pretty low to be significant concern to typical applications.

Now that we have a rough idea on how performance KPIs change under DOCSIS channel conditions, let us focus on specific impact of upstream traffic – both in-home and on the shared DOCSIS channel – to video conferencing applications. All experiments described further were conduction on the “CM under test” which was using the upstream Tmax of 5Mbps – as it indicated the “tipping” point where impact on latencies and jitter can easily be observed.

## 4.2. Understanding Video Conferencing App Operating Points


All three video conferencing apps studied were adaptive to changing network conditions – available throughput, packet latency and jitter as well as packet drops. Each application had a preferred sweet spot for optimal operation, especially for video streaming bandwidth. Table 2 shows typical operating bandwidth range for the video stream from these apps as well as where they tend to “converge” if they have sufficient bandwidth.

**Table 2 – Conferencing App operating “sweet spots”**


| App # | Operating video throughput range | Optimal Target throughput |
|-------|----------------------------------|---------------------------|
| App 1 | 200Kbps – 1.2Mbps                | 700-800 Kbps              |
| App 2 | 500Kbps – 2 Mbps                 | 1-1.25 Mbps               |
| App 3 | 1Mbps – 2.5 Mbps                 | 2.25-2.5 Mbps             |

Conferencing app behavior for media traffic bandwidth with and without network congestion is shown in Figure 6, Figure 7 and Figure 8. It can be observed that some Apps like App 1 are conservative in bandwidth usage while others like App 3 tend to be aggressive in maintaining good video quality pushing continuously to use bandwidth above 2Mbps against traffic congestion. Others like App 2 follow a more step-wise increase approach not unlike several standard adaptive streaming video protocols.


It was observed that at least 500Kbps was required to have minimal acceptable video quality in terms of frame rate and resolution and ideally that close to 1Mbps was required for good quality video streaming. Below 500Mbps, some apps like App 1 continued to show video streams at very low frame rates or video resolution, while other apps moved to “audio only” mode often against the bursty network load.



**Figure 6 – App 1 without (left) and with (right) heavy network load**



**Figure 7 – App 2 without (left) and with (right) heavy network load**




**Figure 8 – App 3 without (left) and with (right) heavy network load**

### 4.3. Speed test & Conferencing App throughput behavior

With multiple conferencing clients behind the test CM, a speed test-like behavior was emulated using TCP tools and the achieved throughput of the speed test app was studied. Figure 9 shows the behavior of such speed tests when 2 conferencing client apps (of the same type) are running in conjunction with a speed test behind the test CM. The baseline speed test shows a typical reduction in achievable throughput with increasing DOCSIS channel congestion (the max upstream is 5Mbps for the test CM).

As seen in the figure below, anomalies in app behavior are observed when network congestion exceeds 75-78% of channel capacity. Interestingly, at lower network congestion levels, the adaptive algorithms were trying to maximize bandwidth use and not let the TCP-based speed test achieve adequate throughput. Once the network gets congested, the algorithms tend to back off and this allows for better speed test results.




**Figure 9 – Speed test along with multiple conference apps**

Looking at the throughput dataset from an alternate perspective, Figure 10 shows how much throughput a single video conferencing app can achieve when a speed test type data is running along with other conferencing sessions behind the test CM. When varied with DOCSIS channel load, this gives an indication of how congested the channel can get before there is deterioration in the conferencing app QoE to the user.

Each App shows a distinct behavior against network congestion.

1. App 1 uses low bandwidth (600-800kbps) when the network is not congested. That is the typical baseline throughput for App 1. However, at around the 75% mark, the bandwidth used starts increasing to around 1Mbps – likely due to packet retransmissions needed to maintain the same video resolution and quality. No impact was seen to the actual video quality of the app until the congestion in the network exceeded 90%.
2. App 2 shows a much more classic adaptive video streaming behavior with a sawtooth-like variation with throughput. While it is not visible from the graph itself, looking at the raw statistics of throughput vs latency, it was evident that the two KPIs worked in tandem with each other. When app latency increased, the video dropped in resolution and quality, allowing less bandwidth to be used.
3. App 3 had the highest bandwidth utilization of the three, typical of the nature of the app as described in Section 4.1. However, it was the most yielding of bandwidth usage, once network congestion hit 75% or higher, dropping to as low as 600Kbps.



**Figure 10 – Speed achieved by a single conference session under CM+channel load**


It should be noted that while having a 1Mbps target for app video throughput is optimal, the app is still very functional at lower throughputs around 500Kbps however with reduced frame rates and/or video quality. All apps were very effective in working with congested networks and working for optimal performance. The tipping point of network congestion where apps tended to react rapidly to changing conditions was around 75-90% of channel capacity.

#### 4.4. Video Conferencing Latency measurements

To measure latency of packets through this congested network, a low bit rate UDP stream was injected from a traffic generator behind the test CM and packet transmission times were measured over the DOCSIS upstream (to the Network Side Interface of the CMTS). Latencies of the data sets were computed and compared, both with and without conferencing apps running, as well as in the presence of the speed-test traffic behind the test CM.


Figure 11 and Figure 12 show mean latencies of conferencing apps under varying upstream DOCSIS channel congestion, with and without additional CM loading with speed test traffic. Both set of graphs have two concurrent conferencing sessions behind the test CM.

Looking at Figure 11, where the home network is not congested, App 3 is showing the highest latency impact. This is to be expected as App 3 had a tendency to take as much as 2.5Mbps per video session if the bandwidth is available. Hence 2 concurrent sessions with additional latency measurement traffic can lead to traffic bursts close to the 5Mbps Upstream Tmax of the service tier of the test CM. App 1 is not impacted much by latency and track very close to the baseline latency (when no conferencing apps are running). App 2 starts getting impacted by latency as network congestion hits 78% and the “sawtooth” behavior mimics the throughput graphs of the same application.




**Figure 11 – Mean Latency for conferencing apps under channel loading ONLY**

In Figure 12, we can now see that once the home network is loaded, latency impact on conferencing apps are much more prominent. Similar to throughput graphs, the inflexion point where apps tend to adapt their reduce their latency is around 78% of network congestion. App 3 being the most bandwidth hungry tends to be most impacted by heavy traffic loads, while App 2 shows a similar behavior to throughput variations with a sawtooth-like fluctuation in latency under heavy congestion.




**Figure 12 – Mean Latency for conferencing apps under CM+channel load**

However, mean latencies only tell part of the story. Jitter or delay variation is key in how these conferencing apps are responding. Most of the packets should be within the latency threshold of 150ms, based on ITU-T G.114 specifications<sup>3</sup> for optimal performance of video conferencing style applications. Let us look at how 95<sup>th</sup> percentile of these latencies to see how the operating points look – to monitor if most of the video packets are maintained under the recommended 150ms threshold.



**Figure 13 – 95<sup>th</sup> percentile latency for conferencing apps under channel loading only**



**Figure 14 – 95<sup>th</sup> percentile latency for conferencing apps under CM+channel load**

Based on Figure 13, with only network loading, App 1 latency is hardly impacted by network loading only. App 2 starts showing increasing latencies when network load exceeds 78% of channel capacity. However, its adaptive algorithm is able to keep latencies under 150ms by reducing video throughput. App 3 shows increasing latencies beyond acceptable threshold when network congestion is around the 78-80% mark.

Figure 14 shows the 95<sup>th</sup> percentile latency of conferencing apps when the TCP speed-test is also running behind the test CM. The apps are much more susceptible to latency impact with this increase load at home, but the inflection point seems to be still around the 78-80% of channel capacity. While increasing in-home traffic impacts App 2 more, App 3 is actually showing improved latencies when the congestion at home increased. One possible explanation for this could be the fact these applications are adaptive in nature. The presence of additional traffic at the CM and algorithms like Active Queue Management could result in better latency metrics.



## 5. Conclusion

Our experiments with various video conferencing applications indicate that these applications adapt well with varying network conditions and traffic patterns. To keep the desired throughput higher than 1Mbps for video quality and less than 150ms of latency for majority of the traffic, our study shows that a DOCSIS upstream channel utilization should not exceed 75-80% of its full operating capacity. Additionally, ensuring optimizations at the cable modem or home gateway level such as Active Queue Management, Buffer Control TLVs and TCP ack suppression will help reduce the impact of a highly congested channel and allow conferencing apps to adapt better to changing network conditions.

Our study focused on a service tier with fairly stringent bandwidth constraints. To mitigate the impacts of home congestion on these applications, it is desirable to move more subscribers to higher tiers. In the medium term, adding any additional available DOCSIS spectrum, leveraging DOCSIS 3.1 technologies like OFDM/OFDMA and performing physical node splits for smaller service group sizes can help mitigate pressure from such bandwidth surges. Long term solutions include moving to mid or high splits, moving to fiber-deep and distributed access architectures including Full Duplex DOCSIS and Extended Spectrum DOCSIS solutions.

## Abbreviations

|          |                                                                   |
|----------|-------------------------------------------------------------------|
| AQM      | Active Queue Management                                           |
| ATDMA    | Advanced Time Division Multiple Access                            |
| CM       | Cable Modem                                                       |
| CMTS     | Cable Modem Termination System                                    |
| COVID-19 | Coronavirus Disease 2019                                          |
| CPE      | Customer Premises Equipment                                       |
| DOCSIS   | Data Over Cable System Interface Specifications                   |
| DS       | Downstream (In the context of DOCSIS transmission – to the home)  |
| KPI      | Key Performance Indicator                                         |
| MSO      | Multiple System Operators (DOCSIS cable operators in our context) |
| OFDM     | Orthogonal Frequency Division Multiplexing                        |
| OFDMA    | Orthogonal Frequency Division Multiple Access                     |
| PC       | Personal Computer                                                 |
| QoE      | Quality of Experience                                             |
| SG       | Service Group                                                     |
| SLA      | Service Level Agreement                                           |
| TCP      | Transmission Control Protocol                                     |
| TLV      | Type-Length-Value                                                 |
| UDP      | User Datagram Protocol                                            |
| US       | Upstream (In the context of DOCSIS transmission – from the home)  |
| VPN      | Virtual Private Network                                           |

## Bibliography & References

1. Jeff Finkelstein, Tom Cloonan, Doug Jones, “MSO Learnings from COVID-19”, SCTE, 2020  
SCTE Cable Tec Expo, 2020.
2. John Ulm, Tom Cloonan, “Managing the Coronavirus Bandwidth Surge”, SCTE, 2020  
SCTE Cable Tec Expo, 2020
3. ITU-T Recommendation G.114, One way transmission time, 05/2003

# Session Overhead Reduction in Adaptive Streaming

A Technical Paper prepared for SCTE•ISBE by

**Alexander Giladi**  
Fellow  
Comcast  
1899 Wynkoop St., Denver CO  
+1 (215) 581-7118  
alex\_giladi@comcast.com

# Table of Contents

| Title                                          | Page Number |
|------------------------------------------------|-------------|
| 1. Introduction.....                           | 3           |
| 2. Session overhead and HTTP compression ..... | 4           |
| 3. Reducing traffic overhead .....             | 5           |
| 3.1. MPD patch.....                            | 6           |
| 3.2. Segment Gap Signaling .....               | 7           |
| 3.3. Efficient multi-DRM signaling .....       | 8           |
| 4. Reducing the number of HTTP requests .....  | 9           |
| 4.1. Asynchronous MPD updates.....             | 9           |
| 4.2. Predictive templates.....                 | 10          |
| 4.3. Timeline extension .....                  | 10          |
| 5. Conclusion.....                             | 11          |
| Abbreviations .....                            | 12          |
| Bibliography & References.....                 | 12          |

## List of Figures

| Title                                                          | Page Number |
|----------------------------------------------------------------|-------------|
| Figure 1: Anatomy of DASH-based streaming application[2] ..... | 4           |
| Figure 2: MPD Patch example .....                              | 7           |

## List of Tables

| Title                                                                                   | Page Number |
|-----------------------------------------------------------------------------------------|-------------|
| Table 1: MPD overhead .....                                                             | 5           |
| Table 2: MPD size (in bytes) with and without HTTP compression .....                    | 5           |
| Table 3: Traffic overhead in megabytes for 1-hr session.....                            | 6           |
| Table 4: MPD size (in bytes) with referencing and HTTP compression .....                | 9           |
| Table 5: Request and traffic overhead of a 1-hr session with asynchronous updates ..... | 11          |

# 1. Introduction

In typical linear adaptive streaming deployments, much thought is given to bitrate optimization, in order to maximize a viewer's quality of experience. The non-negligible overhead of the manifest traffic is, however, typically overlooked. A manifest contains information essential for streaming a video asset, identifying the contents of the stream and the location of where constituent components, like URLs, can be found. A manifest which is frequently refreshed can significantly impact bandwidth consumption and number of requests made to Content Delivery Networks (CDNs). This is true for both Apple® HLS and MPEG DASH streaming systems. While the “manifest bloat” problem is endemic for both systems, this paper concentrates on MPEG DASH and DASH-specific tools. In the case of MPEG DASH, there are several major sources contributing to manifest growth, such as the number of included content periods and the sheer volume of DRM license information. Frequent requests further compound the burden of what is essentially “manifest bloat.” Altogether the manifest overhead can easily reach 250Kbps and go past 2Mbps in some pathologic cases. These numbers are well beyond the typical low-rate video bitrate.

There are several different approaches to minimizing the manifest overhead. Some are as simple as using HTTP lossless compression algorithms, such as gzip, or the more exotic Brotli [11]. Current DASH practices, such as predictive templates, events, and the timeline extension process, provide an orthogonal approach. Operational experience with large-scale, DASH-based linear programming resulted in a set of new manifest reduction tools that were consequently added into the recent version of MPEG DASH. For example, a recently introduced MPD patching mechanism dramatically reduces the manifest overhead by only sending updates when possible. Several other additions reduce the size of the DASH manifest, by reducing the “bloat” due to license acquisition information in multi-DRM content and segment losses.

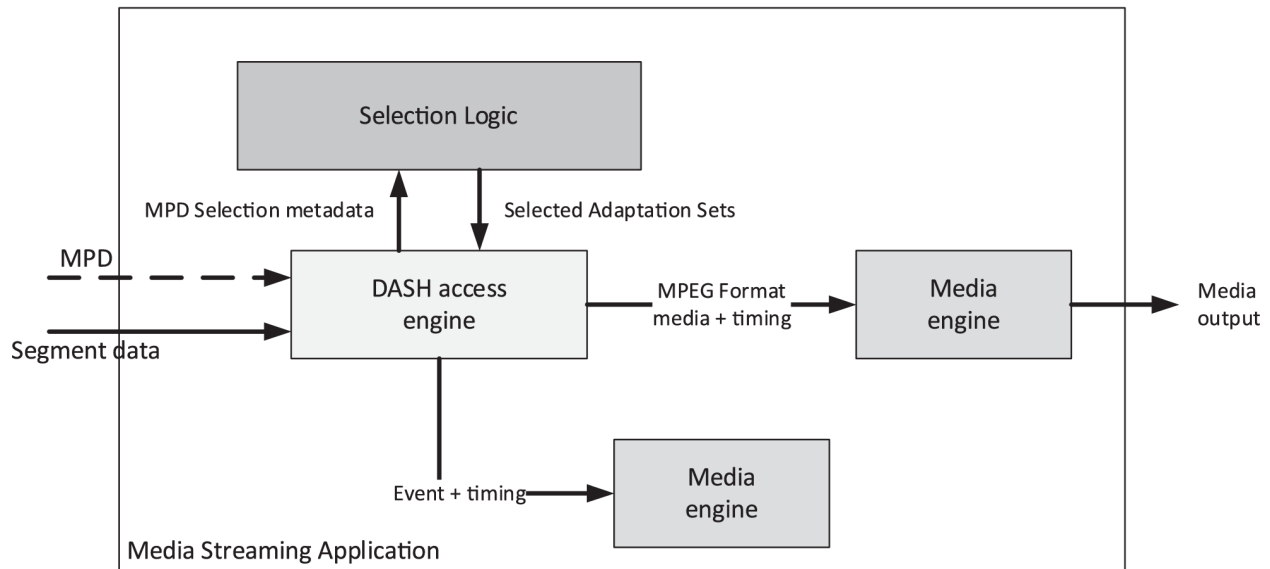
Adaptive streaming over HTTP emerged as the mainstream method of delivering video IP-based networks roughly a decade ago as a response to the challenges of network unpredictability and device heterogeneity. Adaptive streaming technology is what powers over-the-top (OTT) services such as Netflix, YouTube, Hulu, Disney+, Peacock, and HBO Max, among others. In aggregate, adaptive streaming techniques reach hundreds of millions of viewers on a dizzying variety of devices.

Adaptive streaming is by no means a new technology. C. Gonkin et. al. [1] wrote the one of the earliest papers on adaptive streaming describing the Real Networks implementation of the approach. The technology became mainstream much later, after the introduction of the first version of Apple HTTP Live Streaming [6] (HLS) in 2009. Several additional streaming systems, such as Microsoft's HTTP Smooth Streaming (HSS) and Adobe's HTTP Dynamic Streaming (HDS) emerged at around the same time and became viable alternatives. Two years later, MPEG Dynamic Adaptive Streaming over HTTP [2] (DASH) became an international standard. Today, HLS and DASH are responsible for the vast majority of adaptive streaming deployments. DASH also made its way into broadcast as a part of Advanced Television System Committee (ATSC) 3.0 and Europe's Hybrid Broadcast Broadband TV (HbbTV).

As opposed to traditional broadcast, cable, and IPTV systems, which push content to receivers, adaptive streaming is a “pull” system, where a streaming client makes an autonomous decision about what to download based on current network conditions and device capabilities. The content is encoded in multiple bitrates referred to as *representations*. Each representation is encoded as series of short *segments* (playable pieces of video, typically 2-10 seconds each). Representations which have the same media content (e.g. video at different bitrates) and have aligned segment boundaries are grouped into *adaptation sets*. The complete asset (e.g. a movie, a pre-recorded or live show) is broken into one or more independent *periods*, which cover a uniform period of time within a presentation. While video on demand (VOD) assets without advertisements consist of a single period, linear channels typically contain multiple periods. An example of such a multi-period asset can be a linear channel where the first period covers

entertainment content, the next three cover three different advertisements, and the fifth covers the continuation of the entertainment content.

HLS is similar. Its manifest is a collection of media playlists listing segment URLs, and a master playlist that references the media playlists and lists their properties. The media playlists are conceptually identical to representations, while the master playlist combines the roles of both period and adaptation sets.



**Figure 1: Anatomy of DASH-based streaming application[2]**

A DASH streaming session starts with downloading a Media Presentation Description (MPD), which is an XML document containing information about media segments, their timing, and the inter-relationships between them. After parsing the MPD, the client selects the representations it sees fit, given its capabilities and network conditions. It starts downloading the media segments from the selected representations. The segments are typically kept in a player buffer, and eventually decoded and rendered. The client also continuously estimates available bandwidth and monitors buffer fullness, and given these it re-evaluates its representation selection.

A. C. Begen et al [8] provide an excellent introduction to the concepts of adaptive streaming. A much later work by A. Bentalb et al. [9] provides an overview of modern rate adaptation techniques.

## 2. Session overhead and HTTP compression

Any streaming session starts with downloading the manifest. In cases of linear content, this manifest is periodically updated. HLS downloads the master playlist at the beginning of the session, and refreshes the media playlists per each segment [6]. While DASH has several tools for avoiding unnecessary MPD requests, many naïve implementations do not use them, and end up implementing an HLS-like client, where new MPD is requested prior to each segment request.

MPDs can be quite “chatty”. Table 1 below lists linear channel streaming bitrate overheads as measured with several production-grade MPDs, and quantifies the corresponding bandwidth overhead associated with 2-sec segment durations. This 2-sec duration implies an MPD request every 2 seconds (i.e. per each segment).. Two-second segment durations are both a common practice and a recommendation (see e.g. S. Lederer [10]).

**Table 1: MPD overhead**

| MPD | MPD Size (bytes) | Bandwidth (kbps) |
|-----|------------------|------------------|
| A   | 61904            | 247.62           |
| B   | 326933           | 1307.73          |
| C   | 425867           | 1703.468         |
| D   | 552219           | 2208.88          |

We are seeing nearly 250Kbps of overhead with a single-period MPD A, and quickly exceed 1 Mbps as the number of periods in the MPD grows. Using the HLS bitrate ladder described in [7], a 1.3 Mbps overhead translates into the difference between 360p and 540p resolutions. This is hardly negligible. Beyond pure traffic, the size of the MPD also translates into parsing time and memory footprint, both functions of number of XML elements.

HTTP/1.1 allows compression of the body of the HTTP response. This is achieved using transfer coding and happens between the endpoints of the protocol, thus it is mostly transparent to the application. Gzip compression is mandatory in HTTP/1.1, while the newer and widely supported Brotli compression [11] provides noticeably better results. Session-related traffic overhead shrinks dramatically – up to 98% – if HTTP compression is used. – this is illustrated in Table 2 . Even with the most efficient Brotli compression, however, the bitrate overhead is still 50 Kbps for an 18-period MPD D, and nearly 250 Kbps with the ubiquitously supported gzip.

**Table 2: MPD size (in bytes) with and without HTTP compression**

| MPD | Uncompressed | Gzip  | Brotli |
|-----|--------------|-------|--------|
| A   | 61904        | 7096  | 5598   |
| B   | 326933       | 37322 | 5589   |
| C   | 425867       | 17266 | 6583   |
| D   | 552219       | 16885 | 10975  |

We can see that application of HTTP compression is essential to make MPD traffic overhead manageable and reduce the start-up time (due to a much faster MPD download).

The HTTP compression approach only addresses the size of the MPD in bytes on the wire and on the CDN. There are additional aspects which depend on the number of XML elements in the MPD: the memory footprint and the MPD parsing time. These need to be addressed using DASH-specific techniques, which are described in the next section.

### 3. Reducing traffic overhead

This section reviews DASH-specific approaches which can be used to reduce the MPD size, number of XML elements, and parsing time. First we discuss MPD patching, a very powerful tool introduced in the latest amendment to MPEG DASH. Patches provide an extremely significant improvement in traffic overhead. Other techniques are needed to remove unneeded XML elements from the actual XML document, and are instrumental in reducing MPD parsing time and its memory footprint.

### 3.1. MPD patch

In the vast majority of cases the change between the previous and the current MPD is minimal. For example, a new segment may have been added in all representations, and the oldest segment may have been removed. This means that the change affects a minimal number of XML elements. Sending only the difference across consecutive MPDs, as updates, is undoubtedly more efficient than downloading full MPDs. This technique can be far more efficient than the generic application of HTTP compression.

The MPD patching framework was first introduced in 3GP-DASH [13] as MPD Delta. The delta format was line oriented, where operands (insert, remove, replace) applied to lines of text. This approach has a major weakness – it implicitly assumes existence of an actual MPD file on the client. This is often not the case, as clients often store the MPD in an in-memory structure, which may be Document Object Model (DOM) or a custom data structure. Secondly, line-oriented syntax is ill-suited for the case where XML document can be regenerated by different entities – whitespace differences can result in patch application errors. Lastly, the MPD Delta syntax ignored version mismatches – e.g. when a modification made to an MPD which differs from the one used for the creation of the delta file, which may render the MPD constructed in the client memory invalid.

The more recent MPEG DASH patching framework takes a slightly different approach: it operates on XML elements and not lines, and addresses elements using XPath as opposed to line numbers. The syntax of the MPD patch is a very restricted subset of the IETF XML patch framework [15]. The restricted syntax creates a system where there typically is only one way of addressing each specific element, and a mandatory check for the client MPD version precedes each patch application. Amendment 1 [3] to the 4<sup>th</sup> edition of MPEG DASH allows explicit requests for MPD patches as an alternative to requesting a full MPD.

MPD patches change the way the client operates: instead of a simple repeated request to an MPD URL, the client alternates between full MPD and patch requests. When the client starts the streaming session, it downloads an MPD. This downloaded MPD provides URLs for both the next MPD and the next MPD patch. If the client decides to download the patch, the patch process will first validate the MPD version, in order to avoid a mismatch. If the download is unsuccessful, or the patch process fails, the client will request the full MPD as an update.

The results of patch application are strikingly better than those achieved by plain HTTP compression: a single patch for the MPDs listed in Table 2 is 944 bytes uncompressed, and only 349 bytes if Brotli compression is applied.

**Table 3: Traffic overhead in megabytes for 1-hr session**

| MPD | Naïve  | Naïve + Brotli | Patch | Patch + Brotli |
|-----|--------|----------------|-------|----------------|
| A   | 106.26 | 9.61           | 1.34  | 0.678          |
| B   | 561.22 | 9.59           | 1.59  | 0.678          |
| C   | 731.05 | 11.30          | 1.10  | 0.630          |
| D   | 947.95 | 18.84          | 1.81  | 0.683          |

Table 3 shows the impact of using patch updates during a 1-hr streaming session. It assumes 2-sec segments and a period being added every 5 minutes. We can see that the effect of patching is more pronounced than the one of the most efficient HTTP compression alone. For example, in case of MPD A Brotli compression resulted in 90.96% reduction in traffic, while patching alone resulted in a 98.74% reduction. Combining both tools is most efficient, as also reduces the first MPD download by the same 90.96%, which results in a shorter start-up time.



Figure 2 below shows an example MPD which adds 7 new segments. As we can see, it contains the MPD identifier and its version information (publication time) in order to ensure the validity of the patching operation. Another method of ensuring the right patch is downloaded is embedding the version on the new (post-patch) MPD in the URL for the next patch. This example is adopted from [3].

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<Patch
 xmlns="urn:mpeg:dash:schema:mpd-patch:2020"
 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
 xsi:schemaLocation="urn:mpeg:dash:schema:mpd-patch:2020 DASH-MPD-
PATCH.xsd"
 mpdId="42"
 originalPublishTime="2020-05-13T05:34:06+00:00"
 publishTime="2020-05-13T05:34:28.601Z">

 <replace sel="/MPD/@publishTime">2020-05-13T05:34:28.601Z</replace>

 <replace sel="/MPD/PatchLocation[0]">
 <PatchLocation ttl="60">live-stream/patch.mpd?publishTime=2020-05-
13T05%3A34%3A28.601Z</PatchLocation>
 </replace>

 <add sel="/MPD/Period[@id='1']/AdaptationSet[@id='1']/
 /SegmentTemplate/SegmentTimeline">
 <S d="360360" r="6" t="5494659049"/>
 </add>

 <add sel="/MPD/Period[@id='1']/AdaptationSet[@id='2']/
 SegmentTemplate/SegmentTimeline">
 <S d="360960" t="5494660288"/>
 <S d="359040"/>
 <S d="360960" r="1"/>
 <S d="359040"/>
 <S d="360960" r="1"/>
 </add>

</Patch>
```

**Figure 2: MPD Patch example**

### 3.2. Segment Gap Signaling

There is no such a thing as 100% reliability in a complex content origination system. Encoders may fail, packagers may fail, networks may fail. As a result, having ideal continuous sequence of segments is hard to achieve over a long enough period of time. There will always be segment gaps – short periods where a segment was not generated by the transcoder.

Many things can go wrong. For example, when the encoder output is a MPEG-2 TS over UDP and a datagram is lost en route to the packager, the segment is lost as well. If the encoder output is a multicast per representation (a very common configuration), occasional packet loss will affect a single representation at a time.

Another example is an encoder failure – if an encoder fails, for some reason, and a redundant encoder is started and starts outputting segments, there may be a gap of one or more segments, between the last segment written to origin by the primary encoder, and the first segment written by the redundant encoder. When different representations are generated by different encoders, and an encoder fails, there is a gap only for a subset of the representations, and there will be a segment in one or more representations.

One of the advantages offered by DASH periods is that they are independent, and it is possible to vary the number or character of representations across periods. The main driver behind this design was advanced advertising. For example, the entertainment content may carry audio in English, French, and Spanish languages as well as an English narration (video description) audio track. Moreover, some languages will be available as both stereo and multichannel audio. An ad inserted into this content may only have English stereo. This ad will be represented as a separate period in DASH; pre-ad and post-ad periods would also be periods on their own.

In case a segment from a representation is missing, many implementations assume that the representation no longer exists in the presentation. This triggers creation of a new period without that “lost” representation. At some point in time the packager will again start generating segments for the “lost” representations, which will, in turn, trigger creation of a new period.

In the author’s experience, occasional segment gaps resulted in MPDs with 100-300 periods, which made them both exceptionally large and non-trivial for a player. An alternative to this would be using the SegmentTimeline element, which allows gaps in presentation time. With that said, the use of SegmentTimeline in the case of per-representation gaps is also inefficient, as it will need to appear in every single representation -- as opposed to the common practice of including them only at the adaptation set level. For example, in the case of an unencrypted, 12-representation HLS bitrate ladder [7], and a 2-min MPD, the MPD size increased nearly ten-fold, from 11KB to 105KB.

All of this can be avoided at a low cost if a missing segment is explicitly identified. Both DASH and HLS recently allowed such signaling. In DASH this is achieved using a FailoverContent element, which indicates time gaps for which the representation has no segments. This translates to just a few of lines, with a single gap (i.e., one or more consecutive missing segment) corresponding to a single XML element. This has very little impact on either the MPD size or the number of XML elements it contains.

### **3.3. Efficient multi-DRM signaling**

The overhead of DRM license inlining traffic can be significant, especially in context of linear channels with multi-period, multi-DRM MPDs and multiple video and audio options. The Common Encryption standard [4] and several DRMs define a way of including license acquisition information in the MPD. This is needed in order to start the license acquisition in parallel with the download of an initialization segment, as opposed to waiting for it and parsing this information out of a box contained in it. This inevitably grows the MPD size.

For example, consider a single-period MPD with stereo and multi-channel adaptation sets for both English and Spanish, video, and trick modes has 6 cenc:pssh elements. The number is multiplied by the number of DRMs – meaning that the same single-period MPD service with 3 DRMs contains 18 cenc:pssh elements, but only 3 of them are unique.

What makes the problem acute is that these elements are fairly large. For example, the 18-period 320Kb MPD B contains 171Kb worth of license acquisition data embedded in its ContentProtection element. The vast majority of this data is redundant.

The recent amendment to the 4<sup>th</sup> edition of MPEG DASH [3] introduced a referencing mechanism for ContentProtection elements. Unique ContentProtection elements (one for each DRM) are placed once, at the MPD level, and given unique IDs. They are then referenced by ContentProtection descriptors at adaptation set level, which results in 3 unique elements in the reference ContentProtection descriptors and 18 one-line dependent ContentProtection elements. This on its own dramatically reduces the MPD size. For example, the 320Kb MPD B is reduced to 86Kb.

Table 4 illustrates the results of referencing in MPDs A, B, and D, which have 3 DRMs. We can see that referencing results in a very significant reduction in uncompressed size of an MPD. As opposed to HTTP compression, this result translates directly into reduction in memory footprint. With that said, the result of applying lossless Brotli compression to an MPD with referencing is near-identical to applying same Brotli compression to the original MPD.

**Table 4: MPD size (in bytes) with referencing and HTTP compression**

MPD	Periods	Original	ContentProtection Referencing		
			Uncompressed	Gzip	Brotli
A	1	61904	18170	6447	5596
B	10	326933	88010	7350	5568
D	18	552219	155650	16128	11051

Referencing can also be used to reduce the size of an MPD patch carrying a new period, as the license acquisition information would be a significant part of the patch, and it typically does not change every period.

## 4. Reducing the number of HTTP requests

A naïve HLS-like implementation also implies that MPD traffic is responsible for a third of CDN requests. The number of CDN cache misses is lower for the MPDs, but the number of actual requests is still quite large – a third of the requests for an asset containing video and audio. The sheer number of HTTP GET requests is taxing the edge caches.

Several tools in the MPEG DASH specification are intended to reduce the number of HTTP requests for linear content.

### 4.1. Asynchronous MPD updates

DASH has an inband event mechanism largely modeled after SCTE 35 cue messages in MPEG-2 TS. DASH inband events are timed “blobs” of metadata embedded in an Event Message (`emsg`) box. This box resides in the beginning of an ISO-BMFF media segment, before the Movie Fragment (`moof`) box. In order to ensure that the event is received regardless of the representation, all representations within an adaptation set carry the same inband events. A client is expected to parse the very beginning of the incoming media segment, and in case it finds an `emsg` box, it is expected to pass it to the application or process itself.

One key event type defined in the DASH specification is the MPD Validity Expiration event. It lets the client know the time at which its MPD is going to expire. If the presence MPD Validity Expiration event is signaled in the MPD, the client does not need refresh the MPD until explicitly instructed by the MPD Validity Expiration embedded in a media segment. This is a very powerful feature when coupled with the features described in the next two sections – predictive templates and timeline extension.

## 4.2. Predictive templates

Templates are one of the major differences between DASH and HLS. In DASH, templates are mandatory for all DASH Live profiles intended for linear channels and events. A template is a string similar to the one used in printf functions in the C programming language. The template string has several predefined variables, two of which, \$Number\$ and \$Time\$, are of particular interest.

\$Number\$ stands for the number of the requested segment. This number is incremented for every segment, and URLs for every segment are derived by inserting the value of \$Number\$ into the template string. This way, given a more or less constant segment duration, there is no need to update the MPD, because future segment URLs can be derived along with an approximation of their availability window. The DASH-IF guidelines [5] allow segment duration to vary within  $\pm 50\%$  of the target segment duration, however they limit the accumulated drift of any number of consecutive segments to the same 50% of a segment duration. For example, if we assume 2-sec segment duration, each segment can be between 1 and 3 seconds, but the cumulative duration of 1000 segments has to be between 1999 and 2001 seconds.

The \$Time\$ variable leverages the precise value of the start time of the segment. It is used with the SegmentTimeline element, which contains one or more S elements. Each S element represents a sequence of one or more consecutive segments of equal duration, and does this using the principle of run-length coding. For example, a single S element describes 42 consecutive 2-sec segments as a run of 42 segments with length (duration) of 2 seconds each. The start time of the first segment of the run, the run, and the common duration of the segments are all indicated in the S element. Note that the durations are precise -- if the 43<sup>rd</sup> segment is even one frame shorter or longer than that, it will be described in a separate S element.

A special run value (i.e., the value of the S@r attribute) of “-1” means “until further notice”. This way there is no need for an MPD update as long as the segments are precisely identical – start time and duration allow calculation of the value of \$Time\$ and hence the derivation of a segment URL for a segment which has not yet been encoded. This has the same effect as the \$Number\$-based approach, but is easier to use and validate, since the times are precise and not an approximation with a  $\pm 50\%$  tolerance.

Templates using the autoincrementing \$Number\$ or \$Time\$ with run of -1 let us predict URLs of future segments. This way no MPD update is required as long as segments are being made available in time – their URL and the time at which they can be requested can be calculated way before these segments are even created. This works fine if there is no change in the MPD beyond adding and removing new segments.

The likelihood of not having any other changes in a linear manifest is nil – for example, due to advanced advertising which adds new periods and new SCTE 35 events. This commonly leads to an HLS-like implementation where MPD request is issued for every segment. However, the MPD Validity Expiration event can be used to trigger an MPD update before a material change in the MPD, such as a new upcoming period. This approach reduces the number of MPD requests to the bare minimum – potentially, once per every period as opposed to once every 2-second segment. For example, an hour long MPD with an entertainment period, followed by an ad period, followed by the next entertainment period, may require just one MPD update (prior to the start of the ad period) as opposed to 1,800 requests for the traditional per-segment polling.

## 4.3. Timeline extension

The SegmentTimeline predictive mode is based on an assumption of frame-identical segment durations. Given US fractional frame rates, this is not always possible, because there is no reasonably short segment

duration to align a 1024-sample AAC frames, 1536-sample E-AC-3 frames, and 29.97 fps video. As a result, audio segments typically have a duration pattern where periodically adding one longer or shorter segment prevents a drift from forming. This reduces the efficiency of \$Time\$-based addressing by limiting run values in audio adaptation sets.

A more efficient mode, commonly referred to as “timeline extension,” is documented in DASH-IF IOP [5] as “MPD and Segment-based Live Service Offering”. Given that each S element provides a time precisely matching the time in the Track Fragment Decode Time (‘tfdt’) box of the media segment, it is possible to predict the URL of the next segment, having parsed the beginning of the current segment. This approach allows a reliance on event-driven, asynchronous MPD updates and results in the same performance as in the \$Number\$-based case, but with much higher precision.

**Table 5: Request and traffic overhead of a 1-hr session with asynchronous updates**

MPD	Naïve		Asynchronous updates		
	Traffic (MB)	GET requests	Uncompressed (MB)	Brotli (MB)	GET requests
A	106.26	1800	0.71	0.064	12
B	561.22	1800	3.74	0.064	12
C	731.05	1800	4.87	0.075	12
D	947.95	1800	6.32	0.126	12

Table 5 shows the effect of using asynchronous MPD updates (without any of the traffic-reducing techniques from section 3) in a scenario we used in Table 3 above: 1-hr session with 5-min periods and 2-sec segments. We further assume that an MPD Validity Expiration event is sent before the start of the period. This scenario results in a 99.33% reduction in HTTP GET requests and shows the best result in traffic reduction. With that said, the numbers below represent a “happy path”, and every missing segment in each currently playing representation will trigger an extra MPD request. DASH allows a “missing content segment” – a segment which contains no media data and only provides segment timing information. In case of timeline extension, such a segment will prevent unneeded MPD requests.

Note that asynchronous updates can be combined with MPD patches to get to even greater efficiencies.

## 5. Conclusion

In this paper we quantified the impact of “manifest bloat” specific to DASH and its MPD traffic and reviewed several methods of reducing it. While application of brotli compression reduces the traffic on its own by at least 90%, we recommend a combination of the proposed measures. Note that both the timeline extension and the

Measures such as ContentProtection referencing and gap signaling reduce the memory footprint and improve parsing performance by eliminating redundant elements in the MPD.

From a pragmatic standpoint, HTTP compression is the ultimate “low hanging fruit,” in that it carries a significant traffic impact at a fairly low cost. Gap signaling and ContentProtection referencing are relatively easy to implement and operate. MPD patches are also a new feature and requires an implementation effort on both the client and the packager side, with that said its benefits are significant enough to justify those efforts.

## Abbreviations

3GP	Third Generation Partnership
AAC	Advanced Audio Coding
ATSC	Advanced Television Systems Committee
CDN	Content Delivery Network
DASH	Dynamic Adaptive Streaming over HTTP
DASH-IF	DASH Industry Forum
DOM	Document Object Model
DRM	Digital Rights Management
FPS	Frames per Second
HbbTV	Hybrid Broadcast Broadband TV
HDS	HTTP Dynamic Streaming (HDS) (Adobe)
HLS	HTTP Live Streaming (Apple)
HTTP	Hypertext Transfer Protocol
ISO-BMFF	ISO Base Media File Format (a.k.a. mp4)
IETF	Internet Engineering Task Force
MPD	Media Presentation Description
MPEG	Moving Pictures Experts Group
URL	Uniform Resource Locator
VOD	Video On Demand
XML	Extensible Markup Language

## Bibliography & References

- [1] G. J. Conklin, G. S. Greenbaum, K. O. Lillevold, A. F. Lippman and Y. A. Reznik, "Video coding for streaming media delivery on the Internet," in *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 11, no. 3, pp. 269-281, March 2001, doi: 10.1109/76.911155.
- [2] ISO/IEC 23009-1:2019, Information technology — Dynamic adaptive streaming over HTTP (DASH) — Part 1: Media presentation description and segment formats, 4<sup>th</sup> edition.
- [3] ISO/IEC SC29 WG11, Text of ISO/IEC 23009-1 4<sup>th</sup> edition Draft Amendment 1, CMAF support, events processing model and other extensions, available online at [https://wg11.sc29.org/doc\\_end\\_user/current\\_document.php?id=74740&id\\_meeting=182](https://wg11.sc29.org/doc_end_user/current_document.php?id=74740&id_meeting=182)
- [4] ISO/IEC 23001-7:2016, Information technology — MPEG systems technologies — Part 7: Common encryption in ISO base media file format files, 3<sup>rd</sup> edition
- [5] DASH-IF Guidelines for Implementation: DASH-IF Interoperability Points, November 2018, available online at <https://dashif.org/docs/DASH-IF-IOP-v4.3.pdf>
- [6] R. Pantos, HTTP Live Streaming 2nd Edition, IETF I-D, available online at <https://datatracker.ietf.org/doc/html/draft-pantos-hls-rfc8216bis-07>
- [7] Apple Inc., "HLS Authoring Specification for Apple Devices", available online at [https://developer.apple.com/documentation/http\\_live\\_streaming/hls\\_authoring\\_specification\\_for\\_apple\\_devices](https://developer.apple.com/documentation/http_live_streaming/hls_authoring_specification_for_apple_devices)
- [8] Ali C. Begen, Tankut Akgul and Mark Baugher, "Watching video over the Web, part 1: streaming protocols," *IEEE Internet Comput.*, vol. 15/2, pp. 54-63, Mar./Apr. 2011.

- [9] A. Bentaleb, B. Taani, A. C. Begen, C. Timmerer and R. Zimmermann, "A Survey on Bitrate Adaptation Schemes for Streaming Media Over HTTP," in *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 562-585, Firstquarter 2019, doi: 10.1109/COMST.2018.2862938.
- [10] S. Lederer, Optimal Adaptive Streaming Formats MPEG-DASH & HLS Segment Length, November 2015, available online at <https://bitmovin.com/mpeg-dash-hls-segment-length/>
- [11] IETF RFC 7932, Brotli Compressed Data Format, July 2016
- [12] IETF RFC 7232, Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests, June 2014
- [13] 3GPP TS 26.247 'Progressive Download and Dynamic Adaptive Streaming over HTTP (3GP-DASH) v. 2.0.0 (Release 10), June 2011.
- [14] Zachary Cava, Scaling Live OTT with DASH: Techniques and Lessons Learned, Mile-High Video 2019, Denver CO, available at [http://mile-high.video/files/mhv2019/pdf/day2/2\\_16\\_Cava.pdf](http://mile-high.video/files/mhv2019/pdf/day2/2_16_Cava.pdf)
- [15] IETF RFC 5261, An Extensible Markup Language (XML) Patch Operations Framework Utilizing XML Path Language (XPath) Selectors, September 2008

# Cable and Mobile Convergence

## A Vision from the Cable Communities Around the World

A Technical Paper prepared for SCTE•ISBE by

**Jennifer Andréoli-Fang, PhD**

Distinguished Technologist

CableLabs

[j.fang@cablelabs.com](mailto:j.fang@cablelabs.com)

**John T. Chapman**

CTO Cable Access, Cisco Fellow

Cisco

[jchapman@cisco.com](mailto:jchapman@cisco.com)

<b>Charter Communications</b>	Ahmed Bencheikh, Group Vice President Praveen Srivastava, Director Vikas Sarawat, Consulting Engineer
<b>Cisco</b>	Ian Campbell, CTO Mobility, Distinguished Engineer Mark Grayson, Distinguished Engineer
<b>Cox Communications</b>	Drew Davis, Executive Director Paul Blaser, Director
<b>Shaw Communications</b>	Damian Poltz, Senior Vice President Dave Morley, Director
<b>Telecom Argentina</b>	Eduardo Panciera, Chief Architect
<b>Vidéotron</b>	Philippe Perron, Director Sylvain Archambault, Director Eric Menu, Architect Géraldine Trouillard, Architect David Lagacé, Architect
<b>Vodafone</b>	Gavin Young, Head of Fixed Access Centre of Excellence Bruno Cornaglia, Distinguished Engineer, Fixed Access Sr Manager



# Table of Contents

Title	Page Number
1. Introduction.....	5
2. Four Stories of Convergence .....	7
2.1. Business Convergence .....	7
2.2. Infrastructure Convergence.....	9
2.2.1. Transport Network.....	9
2.2.2. Open Source Code .....	9
2.2.3. Data Center.....	10
2.3. Service Convergence.....	10
2.4. Architectural Convergence.....	11
3. Traffic Engineering .....	12
4. A Survey of Mobile Deployment Plans by MSOs Around the World.....	14
4.1. Charter Communications .....	14
4.1.1. Spectrum Mobile Data Offload.....	14
4.1.2. Cellular Wireless – A New and Exciting Frontier .....	14
4.1.3. Possible Wireless Deployment Scenarios .....	15
4.1.4. Backhaul .....	19
4.1.5. Dual-SIM .....	22
4.1.6. Virtualized RAN Architecture .....	24
4.1.7. Fixed-Mobile Convergence .....	25
4.1.8. Automation .....	25
4.2. Cox Communications .....	27
4.3. Shaw Communications .....	28
4.3.1. Inter-Operator Handover.....	30
4.3.2. Access Network Convergence.....	32
4.3.3. Wi-Fi Interoperability .....	33
4.4. Vidéotron.....	35
4.4.1. Who is Vidéotron?.....	35
4.4.2. 5G Opens a Very Wide Range of Applications.....	35
4.4.3. Wireless Access Convergence – Building on Wi-Fi Assets.....	36
4.4.4. Network Convergence – All-in-One Fiber Network.....	38
4.4.5. Virtualization – A Core Ready for All Opportunities .....	40
4.5. Telecom Argentina .....	42
4.5.1. Introduction .....	42
4.5.2. Convergence Access Network and Convergence Access Infrastructure .....	44
4.5.3. Fixed Access Network .....	45
4.5.4. Small Cell Backhaul over DOCSIS .....	46
4.5.5. Evolution Towards 5G.....	48
4.5.6. Convergence Access Infrastructure.....	51
4.5.7. Virtualization and Cloudification.....	52
4.6. Vodafone.....	54
4.6.1. Background and Context .....	54
4.6.2. Motivation for Convergence .....	55
4.6.3. Convergence Use Case Examples and Network Scenarios .....	57
4.6.4. Summary.....	68
5. Convergence – Technologies .....	69
5.1. DOCSIS Technology.....	69
5.2. Low Latency Xhaul (LLX) .....	69
5.3. Synchronization and Timing.....	71
5.4. DOCSIS DAA for Mobile People .....	72
5.5. DOCSIS for Mobile Xhaul .....	73
5.6. Common Quality of Service (QoS) Framework.....	74

5.7.	A 5GC View of Convergence .....	75
5.8.	Managing the RAN with YANG .....	77
6.	Convergence – A Vision of What is to Come .....	79
6.1.	Integrated and Converged HFC, DOCSIS and Mobile Network .....	79
6.2.	Converged Transport of Mobile Xhaul over DOCSIS .....	79
6.3.	Converged Transport with Common CIN .....	80
6.4.	Common Cloud Platform .....	80
6.5.	Virtual and Cloud Native Functions .....	80
7.	Conclusion .....	82
	Abbreviations .....	84
	Bibliography & References .....	87

## List of Figures

Title	Page Number
Figure 1 – Cable MVNO Customer Adds .....	5
Figure 1 – Canadian Service Provider Market .....	6
Figure 2 – The Four Stories of Convergence .....	7
Figure 3 – Subscriber Convergence .....	8
Figure 4 – Architectural Convergence .....	11
Figure 5 – Small Cell Traffic Engineering .....	12
Figure 6 – An Example CBRS Strand Mount Deployment with Overlapping Macrocell .....	15
Figure 7 – Tri-Star Configuration Product .....	16
Figure 8 – Attached Mount Mini-Macro .....	17
Figure 9 – Example Cat-A CBSD .....	17
Figure 10 – Residential Femtocell Deployment .....	18
Figure 11 – Residential Small Cell Gateway Options .....	19
Figure 12 – How a DSDS Device Moves Across Networks .....	23
Figure 13 – vRAN Architecture – Backhaul .....	24
Figure 14 – vRAN Architecture – Midhaul .....	24
Figure 15 – vRAN Architecture – Fronthaul .....	25
Figure 16 – Charter Wireless Automation Framework .....	26
Figure 17 – Wireline and Wireless Convergence Roadmap .....	30
Figure 18 – Home Routed (HR) Network Architecture .....	31
Figure 19 – Converged Access Network Architecture .....	32
Figure 20 – 5G Diversity of Services and Applications .....	36
Figure 21 – Wireless Access Convergence .....	38
Figure 22 – Fiber Network Convergence with DWDM Ethernet Aggregation .....	39
Figure 23 – Fiber Network Convergence with High-Capacity Coherent Ring .....	40
Figure 24 – Convergence and Virtualization .....	41
Figure 25 – Telecom Access Network .....	42
Figure 26 – Telecom Access Technologies .....	43
Figure 27 – Telecom Access Network Prior to Convergence .....	44
Figure 28 – Access Convergence .....	45

Figure 29 – 4G Small Cells Backhaul over DOCSIS Architecture .....	47
Figure 30 – LTE outdoor microcells Throughput.....	47
Figure 31 – 3GPP Deployment Options.....	49
Figure 32 – Possible Evolution Toward 5G.....	49
Figure 33 – Convergent 5G Core.....	50
Figure 34 – Access Transport Technologies .....	51
Figure 35 – Virtual or Cloud Access Services.....	53
Figure 36 – Vodafone's European Presence .....	55
Figure 37 – Infrastructure Convergence .....	56
Figure 38 – Vodafone Four Layers of Convergence.....	57
Figure 39 – Convergence Use Cases.....	58
Figure 40 – Vodafone TV .....	58
Figure 41 – Local Streaming.....	59
Figure 42 – Smart Home.....	59
Figure 43 – Hybrid Access.....	60
Figure 44 – Always-on Service .....	61
Figure 45 – Fixed Wireless Access.....	61
Figure 46 – Converged Access with DWDM PON.....	62
Figure 47 – Network Timing with DOCSIS and PON.....	63
Figure 48 – Unified Fiber Access/Aggregation Network .....	63
Figure 49 – Wireline Support in 5G.....	64
Figure 50 – Wireline Connection to 5GC with AGF .....	65
Figure 51 – Voice over Wi-Fi.....	66
Figure 52 – Quality Attenuation technique.....	66
Figure 53 – USP TR-369 in Deployment .....	67
Figure 54 – DOCSIS GW with TR-369 .....	68
Figure 55 – Scheduler Pipelining with Bandwidth Report (BWR).....	70
Figure 56 – 1588 Timing over DOCSIS Network with DTP .....	71
Figure 57 – Mobile RAN Splits with DOCSIS DAA .....	72
Figure 58 – Mobile Xhaul over DOCSIS Architecture .....	74
Figure 59 – Common Quality of Service Framework for MBH over DOCSIS.....	74
Figure 60 – 5G Convergence.....	75
Figure 61 – Platform, Policy, and System Convergence .....	76
Figure 62 – Common Cloud Native Platform .....	77
Figure 63 – DOCSIS-Mobile Network Convergence .....	79


## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Operator Convergence .....	8
Table 2 – RF Cell Radius Comparison .....	13
Table 3 – Small Cells per Fiber Node .....	13
Table 4 – Timing and Sync Requirements of LTE and 5G .....	20
Table 5 – Options to Provide Timing and Sync.....	21

Table 6 – Comparison of Three Dual-SIM Technologies.....	22
Table 7 – DOCSIS Capabilities.....	69

## 1. Introduction

After tremendous growth and success in video, data, and voice services, cellular wireless is the next frontier for cable. Quarter over quarter, the US cable mobile virtual network operator (MVNO) business continues to see mobile subscriber growth. As of the Q2, 2020, less than three years after the launch of the first MVNO by Comcast, the three US cable MVNOs combined have amassed 4.2 million customers [1]. This is shown in Figure 1.



**Figure 1 – Cable MVNO Customer Adds**


While Comcast and Charter MVNO utilize Verizon as the mobile network operator (MNO), Altice USA has an agreement with Sprint (now T-Mobile). Recently, Cox Communication also demonstrated an interest in starting an MVNO.

The momentum is there. Consistently, executive leadership at cable companies has shown strong support for and interest in growing the wireless business.

The mobile customer growth is not restricted to the cable operators in the US. Cable operators in the neighboring north – Canada – are also reporting impressive number and innovative business model.

In contrast to the US, virtually all of Canada’s largest cable and telco operators offer both wireline and mobile services on their own infrastructure. Rogers, Canada’s largest cable and mobile operator, has been offering mobile services since 1985, with Vidéotron launching its wireless services in 2010, Shaw acquiring Wind Mobile in 2015, and Cogeco aiming to enter the Canadian wireless market through a Hybrid Mobile

Network Operator (HMNO) model. The Canadian market also faces strong competition from Canada's large incumbent Telco operators, which have invested heavily in fiber to the home, connecting more than 60% of their broadband homes directly to fiber, and leveraging a robust RAN sharing agreement to minimize their infrastructure costs.



**Figure 2 – Canadian Service Provider Market**

In the South America market, Telecom is a leader in the Communication Service Provider (CSP) industry in Argentina. It is a merger of two companies – the former Telecom, incumbent telco, and former Cablevision, a cable operator. The merger of companies was approved in 2018, and the name TELECOM was kept as the name of the new company.

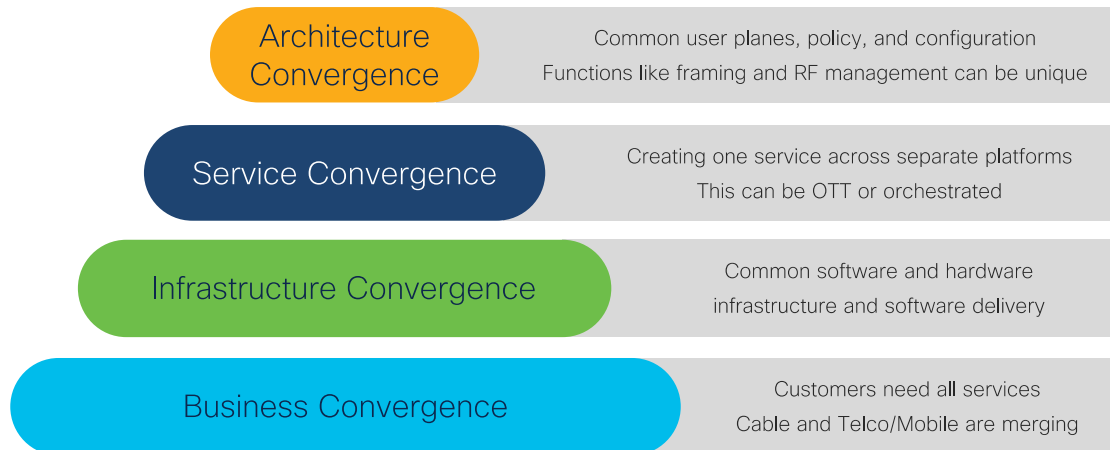
Telecom is the first CSP in Argentina to provide quad-play services: it serves 4 million fixed broadband subscribers, 18.8 million mobile subscribers, 3.5 million TV subscribers and 3 million fixed voice subscribers. It also provides business services. Telecom Argentina is a connectivity solutions and entertainment company transforming the digital experience of its customers, providing them a secure, flexible and dynamic service on all of their devices, with high speed mobile and fixed connections, and a live and on-demand contents platform which includes series, films, gaming, music and TV shows. It is also present in Paraguay, providing mobile service, and in Uruguay, with pay TV.

In Europe, the transformation of cable company to mobile company is mostly done. One of the biggest cable companies in Europe, and the world for that matter, is Vodafone. Vodafone acquired a variety of cable properties across Europe and now operates cable franchises in Germany, Spain, Czech Republic, Hungary, Romania, Albania.

The mobile journey of the cable companies will become an amazing success story. This white paper is intended to grab a snapshot of that story as it is coming together today. The paper will start with a framework for looking at different types of convergence which is also different ways of investing capital for different outcomes. The paper will then showcase six different cable operators with mobile plans, each of whom have common goals and technologies, but a unique point of view. Once the goals and objectives are understood, the paper will highlight technologies that are common to the solutions. These are some of the important technologies that we want the industry to focus on.

## 2. Four Stories of Convergence

There is a lot of buzz round convergence. What is convergence? Is it something in a 3GPP specification? Is it a product? Is it something in your architecture? Is it your customer's experience? Is it your experience? In this section, we would like to describe four stories of convergence. These are shown in Figure 3.



**Figure 3 – The Four Stories of Convergence**

### 2.1. Business Convergence

First things first. It does not make sense to converge technologies and architectures if there is not a business reason. As a stark example, it does not make sense to merge a DOCSIS and 5G system if your company is a pure mobile / telco company that does not have cable properties. Conversely, if you are just a cable operator with no mobile operations, does a 5G driven CMTS add profit to the bottom line? It doesn't.


So, the first step of business convergence is subscriber convergence. The same subscriber needs to pay for both cable broadband service and mobile service. This has already happened across the world. As you are reading this paper, you own a smart phone, a laptop, and have some kind of broadband service at home with either a legacy video service or some type of over the top video service, as does everyone in your family and so do your friends. Subscriber convergence is a done deal.

The second step in business convergence is convergence of cable and mobile operators. As shown in Table 1, this is currently happening around the world. Today's cable operators are tomorrow's mobile operators.

**Table 1 – Operator Convergence**

Region	Cable Operators	Status
USA	GCI	MNO
	Comcast, Charter, Cox, Altice	MVNO
Canada	Shaw, Rogers, Videotron, Eastlink	MNO
Latin America	Telecom Argentina, Claro Brazil, Liberty Latin America	MNO
Europe	Telenet (Belgium), SFR (France), Vodafone (Germany, Spain, Czech Republic, Hungary, Romania, Albania), NOS (Portugal), Telenor (Norway), TDC Net (Denmark), VodafoneZiggo (Netherlands), Telia (Norway/Finland), Tele2 (Sweden/Baltic), Virgin Media+Telefonica/O2 (UK)	MNO
	Virgin Media (UK), UPC Poland, UPC Switzerland	MVNO

Now, it would make sense that if you were to get both your cable broadband service and your mobile service from the same operator, that you would want to get a discount? What is that impact to the operator?



**Figure 4 – Subscriber Convergence**

Let's say you pay \$100 a month for broadband and a \$100 a month for mobile. If you combined your services with one operator, getting both services for \$150 would seem like a reasonable deal. This is shown in Figure 4. Now, as an operator who now owns both the mobile property and the cable property, this means less overall revenue for the same customer base. That means that either one of those operations – cable or mobile, needs to reduce its operating expenses by 50%, or both operations have to reduce their operations by 25%.

An important aspect of converging cable and mobile businesses together is organizational convergence. Otherwise, you just have two companies within one company, sometimes each with their own network. The challenge and importance of this is not to be underestimated and it often requires some sort of cultural convergence. There can no longer be a cable vs mobile competition within the new company. Everyone

must work together to make their subscribers happy with their services at the best cost operating model for the new company.

Subscriber convergence has already happened. Business convergence is happening now. Technology convergence is next [2]. And the outcome that overall convergence should strive for should be at least a 25% reduction in operating costs, if not more.

## **2.2. Infrastructure Convergence**

If you could wear the same set of clothes every day and for everything you do, you would save a lot of money in clothing. Well, that may be impossible for most people, but what about converging the network that hosts a DOCSIS and mobile service? Could that save some money? There are many opportunities for infrastructure convergence. Here are a few examples.

### **2.2.1. Transport Network**

Macrocells and small cells are typically connected using dedicated fiber. For new small cell locations, this might necessitate pulling fiber to a new location along with some copper for powering. This may prove cost prohibitive if a large number of fibers have to be pulled just to support a new service. Alternatively, the cable HFC plant is already composed of fiber and copper, each of which provide an opportunity for transport convergence.

The fiber portion of the HFC network currently uses an “analog” transport that is composed of modulated wavelengths. As the HFC plant gets upgraded to the new Distributed Access Architecture (DAA), that fiber will be converted to “digital” which refers to Ethernet over a wavelength. The analog lambdas are sensitive to noise and cannot coexist with digital lambdas.

As discussed in this paper, the HFC plant with DOCSIS as the transport layer is a great choice for mobile xhaul. So, just using HFC/DOCSIS as backhaul for mobile, even though DOCSIS and mobile might be two different services, is an example of transport convergence that has the potential to save a considerable amount of money by not having to deploy more fiber. By connecting the small cell to the coax side of the HFC plant, the fiber can be analog or digital.

While the HFC plant is still analog, analog services and digital services can use the same plant but must be on separate fibers. When the HFC plant is upgraded to DAA and the fiber becomes digital, then different services can be connected to the fiber, either through dense wave division multiplexing (DWDM), with one service per wavelength, or with an Ethernet switch located in the field that aggregates services on separate 10 or 25 Gbps links and then backhauls them all together on one 100 Gbps to 400 Gbps wavelength.

### **2.2.2. Open Source Code**

Almost all products these days use open source. One of the nice things about open source software is that there are a lot of creative solutions and interesting choices. Open source can save considerable development time. One of the downsides of open source is support. Open source code goes in and out of style. If your system adopts some popular open source that suddenly becomes unpopular, you will lose the support of the community. This means it will be harder to add features and get bug fixes.

Another challenge for open source is that there are security holes and sometimes hundreds of them. Each release of open source needs to be properly vetted to close those security holes. And then there is personnel training. Each piece of software needs to be installed, supported, and upgraded by someone.



Wouldn't it make sense to use the same revision of open source across multiple systems? Would it not make sense to have a uniform code security policy? Wouldn't it make sense to use the same Kubernetes deployer for microservices?

To be fair, much of this open source is very transparent to the operator and is packaged internally by the vendor. There could be separate vendors for cable and mobile systems. Still, tracking open source issues is a real issue to consider.

### **2.2.3. Data Center**

Software used to run in dedicated hardware boxes. Then in the mobile world, software was virtualized and placed onto servers. Now the software is being rewritten using cloud native technologies with microservices and containers intended for an SP edge (data center) environment.

This is cloud and edge ready software. Unfortunately, not all data centers environments are the same and cloud native technologies have been evolving at such a rapid pace recently, that it is currently an operational challenge to completely mix different applications from different vendors into a common cloud or data center deployment. But, it will happen, and this will help achieve a goal of operational simplicity where servers are installed in server farms and software does not much care where it lands.

Infrastructure convergence is happening and is a natural consequence of getting a product to market in an efficient manner. There is more work to be done, but it is a well understood problem that can be worked and optimized.

## **2.3. Service Convergence**

A service is something a customer receives, experiences, and pays money for. It has a meaning and a feeling. The challenge is to identify those services that cross the mobile/cable boundary that can generate additional revenue, reduce operating expense, or both.

A converged service would imply that it should not matter where you are, or if you are connected by mobile, Wi-Fi, DOCSIS or PON -- you get the same service and experience.

Some examples of converged services that come up are the following:

**Parental Controls:** A parent can assign parental controls to a child's device and those controls work whether the device is on the cable broadband network or on the mobile network

**Service Class Roaming:** If you are a premium subscriber with a higher bit rate, you get the same service at your neighbors or the local coffee shop, even if those locations have lower service

**Shared credentials on Mobile/WiFi:** WiFi is often the extension of a DOCSIS network. This convergence would allow a single sign on and may even allow the network to direct which path to take based upon local mobile or cable network loading.

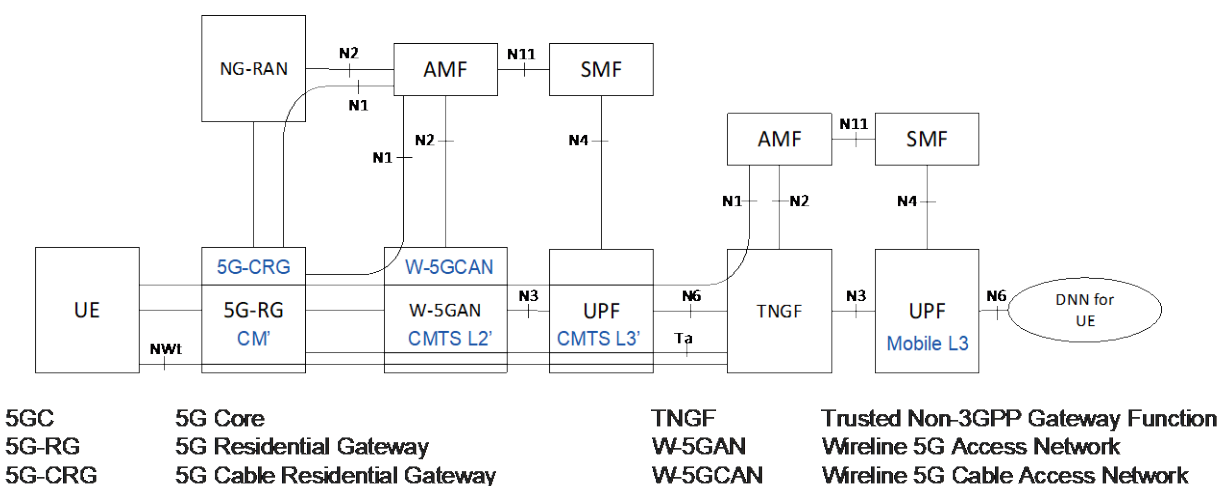
**High Availability:** If one of your transports such as cable access goes down, your services would transfer to the mobile access. This could occur transparently in the network or with some kind of blended service environment in the home.

It should be noted that mobile service today is per device – typically per cell phone and this per user. Cable broadband service is per home and the billing entity is per home which is per cable modem. There is a difference in device identification technologies. Cable modems do not use SIMs. Cell phones do.

One example of distinct separate networks but with a converged service is video conferencing where the video stream is sent over the broadband Internet, but the voice is sent over mobile as the mobile voice connection seems more reliable than the Internet connection. This convergence is simply achieved by the video using an IP address and the audio using a telephone number.

The key question is if the networks need to be converged to achieve these goals, or can dissimilar networks like mobile and DOCSIS just share common policy?

## 2.4. Architectural Convergence



**Figure 5 – Architectural Convergence**

Architectural convergence is where the cable and mobile networks are truly converged. They share a common user plane and a common control plane. This convergence allows for different messaging over the RF interface. The classic example is the wireless and wireline convergence (WWC) work done in 3GPP [3][2]. This is shown in Figure 5, annotated with DOCSIS and mobile functionality. The groundwork for this architecture is partly based in joint CableLabs-Cisco whitepaper at a previous SCTE Expo [4].

Does this architectural convergence add simplicity or complexity? The DOCSIS system works fine and has worked fine for 25 years. Does drastically changing a 25-year old system add value or just add complexity?

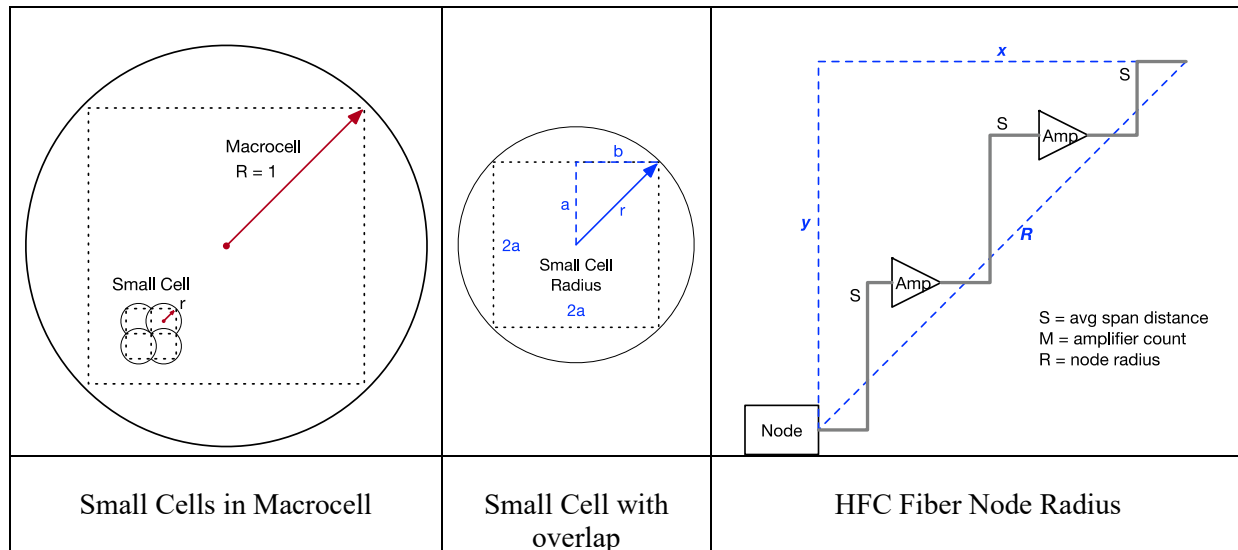
What about software upgrades? Do the mobile and DOCSIS systems now have to be upgraded at the same time since that may be how they were tested? How does that impact the service velocity, which is the ability to roll out new services and bug fixes?

Does it reduce OPEX and CAPEX costs or does it add to those costs? Can you still mix and match vendors when the systems are highly interconnected? That may impact the operator's ability to buy best of breed. Can the entire DOCSIS CM provisioning system be replaced or augmented with a 5G provisioning system? Is there enough financial motivation for vendors to implement these changes?

These are important questions that are not fully answered yet, even though the early architectural pieces have been put into place. One thing is certain though, and that the work does not begin here, it ends here. Before architectural convergence takes place, there has to be business, infrastructure, and service convergence first.

### 3. Traffic Engineering

As the cable network adapts to the connectivity of small cells to fiber or HFC, it is instructive to know how many small cells are required to replace or supplement a macrocell installation. Those numbers are calculated in [5] and briefly summarized here.



**Figure 6 – Small Cell Traffic Engineering**

An ideal overlap of one a small cell radius to another is equivalent to placing a square in a circle and seeing how many small squares (for small cells) fit into a large square or circle (for macrocell). A similar calculation can be done by calculating the radius of the coax segment of an HFC plant. This is shown in Figure 6.

The number of small cells required depends upon the ratio of the service radius and is given by the formulas:

$$\#SC \text{ per } MC = \pi/2 (R/r)^2$$

$$\#SC \text{ per } FN = \pi/4 (S/r)^2 (M + 1)^2$$

where:

$R$  = larger radius of macrocell or node

$r$  = smaller radius of small cell

$S$  = average coax span between actives

$M$  = number of amplifiers in an  $N+M$  cascade

The results of these formulas are shown in Table 2 and Table 3. The relative radius in Table 2 assume same power and same height. In reality, the small cells will be lower power and lower height, plus there will be RF blockage from tree, hill, and walls. As a result, actual deployed results may be higher. Likewise, if less than 100% coverage is needed, the numbers scale back down.

**Table 2 – RF Cell Radius Comparison**

Band	Service	Cell Type	Relative Radius	# Radios
700 MHz	LTE	MC	1.0	1
3.5 GHz	CBRS	SC	0.09	125 to 200
28 GHz	mmWave	SC	0.003	110,000 to 175,000

The first conclusion is that it can the number of radios for CBRS can be 100x that of an LTE macrocell, and for mmWave, the number of radios could be 100,000x that of an LTE macrocell. This dramatically impacts deployment economics and makes the existing HFC plant an interesting choice for mobile backhaul.

**Table 3 – Small Cells per Fiber Node**

		Small cell radius r:		500	1000	2000
M Amps	Total Span	Node Radius	Avg Span	# radios	# radios	# radios
0	1100	778	1100	4	1	1
1	2000	1414	1000	13	3	1
2	2850	2015	950	26	6	2
3	3700	2616	925	43	11	3
4	4600	3253	920	66	17	4
5	5460	3861	910	94	23	6
6	6300	4455	900	125	31	8

A second conclusion is that there is not just one small cell per fiber node on an HFC plant. It depends on the size of the coax segment the HFC plant. So, an N+0 plant may need one to four small cells, where as a N+5 plant may need 25 to 100.

## **4. A Survey of Mobile Deployment Plans by MSOs Around the World**

Now that there is a framework and traffic engineering that show the need to connect small cells in an economical way, let's look at the goals and achievements of six prominent cable operators who are also MNO or MVNO as well.

### **4.1. Charter Communications**

Since the merger earlier in the decade, Charter has been bullish on wireless and continues to invest significantly in research and development. Charter participation and leadership in industry forums and investments in startup ecosystem are a few obvious evidences of the long-term interest and strategy.

#### **4.1.1. *Spectrum Mobile Data Offload***

In the second half of 2018, Charter launched a mobile service (Spectrum Mobile) via a Mobile Virtual Network Operator (MVNO) agreement with Verizon Wireless. All Spectrum Mobile voice and SMS traffic is transported over Verizon Wireless' cellular network. Spectrum Mobile data traffic is either transported over Verizon's cellular network or over Wi-Fi networks.

One of the core technologies enabling cable MVNO is wireless data offload. Charter has an extensive indoor and outdoor Wi-Fi network, which, in addition to providing broadband services, is used to offload the mobile data traffic away from Verizon's network. The contribution of Wi-Fi data-offload in making Spectrum Mobile a fast-growing and economical business can't be overstated.

However, there are limitations on where and how much Wi-Fi can be effectively used to offload cellular traffic. There are some well-known challenges with Wi-Fi data offload. For example,

- Some Spectrum Mobile customers manually switch off the Wi-Fi interface on the End-User Devices
- In congested networks, it isn't easy to manage the Quality of Experience (QoE) on Wi-Fi networks

As a result of these challenges with Wi-Fi, a decent size of the "off loadable" traffic is not offloaded from the Verizon network.

Operators are evaluating multiple solutions and tools to improve the offload numbers. One solution of particular interest is CBRS small cell, which could be used to complement Wi-Fi networks. Not only the CBRS 4G LTE networks could help with the offload, but also offer better QoS and mobility than incumbent Wi-Fi networks.

The CBRS RAN could target a range of indoor and outdoor deployment scenarios, such as outdoor, strand installations covering high demand areas and indoor, small and medium business (SMB), enterprise, residential femtocell deployments. These deployment scenarios are described in more detail later.

#### **4.1.2. *Cellular Wireless – A New and Exciting Frontier***

The deployment of new wireless technologies from scratch is not without challenges. Cable companies have years of experience designing, deploying and managing Wi-Fi based networks both outdoors and indoors.


The 3GPP wireless technologies have unique requirements, which require attention of the cable operators and vendors alike. We talk about these unique requirements in the following sections to draw attention of the cable ecosystem and work together to conquer the wireless frontier at scale.

### 4.1.3. Possible Wireless Deployment Scenarios

#### 4.1.3.1. Fixed Wireless Access (FWA)

FWA can target rural areas to extend the broadband service to unserved or underserved areas, and expand the footprint of the HFC network economically. FWA requires a transceiver outside the home. This use case can be realized by installing RAN equipment onto telecommunication towers and/or other suitable mounting locations, such as water towers.

#### 4.1.3.2. Strand Mount (Aerial)



**Figure 7 – An Example CBRS Strand Mount Deployment with Overlapping Macrocell**


This is a one box solution that is mounted on the strand, where wireless radio is integrated with a DOCSIS 3.1 cable modem (CM). CBRS Category-B device types are planned with a quasi-omni and dual sector design.

Virtual RAN Deployment Model: operators are exploring technologies with split option 2 (shown later in Figure 58) for its vRAN based 5G CBRS wireless network deployment over DOCSIS because of its data transmission latency advantages. Since the DOCSIS network provides advanced latency reduction features such as the Bandwidth Report in the Low Latency Xhaul (LLX) technology [6], Charter may leverage such a feature to further ease CBRS wireless network deployment on cable strand. We will keep monitoring the latest advancements in vRAN field, and latest releases of 3GPP standards, together with O-RAN Alliance initiatives to fine tune our vRAN network deployment strategies.

Strand-mount appears to be the most cost-effective solution where aerial cable strand lines exist. The size and power of strand-mounted solutions are lower than attached mount (explained in a later section) CBRS device (CBSD), but their biggest limitation is power consumption from the HFC plant power supply. Another potential limitation of a strand unit is the mounting orientation – it has to be always along the strand and thus hotspot-targeted deployment in this case might be challenging. To mitigate this, engineers are evaluating quasi-omni strand design that has dual sectors with two sets of antennas covering NE and SW directions. Their height is always 18 feet and typically comes with 2 transmitters 2 receivers (2T2R)

multiple-input multiple-output (MIMO) capability. Utilizing existing assets to mount, connect to, maintain and operate may make these the most cost-efficient deployment type for outdoor Category-B CBSDs.

#### **4.1.3.3. SMB Inside-Out**



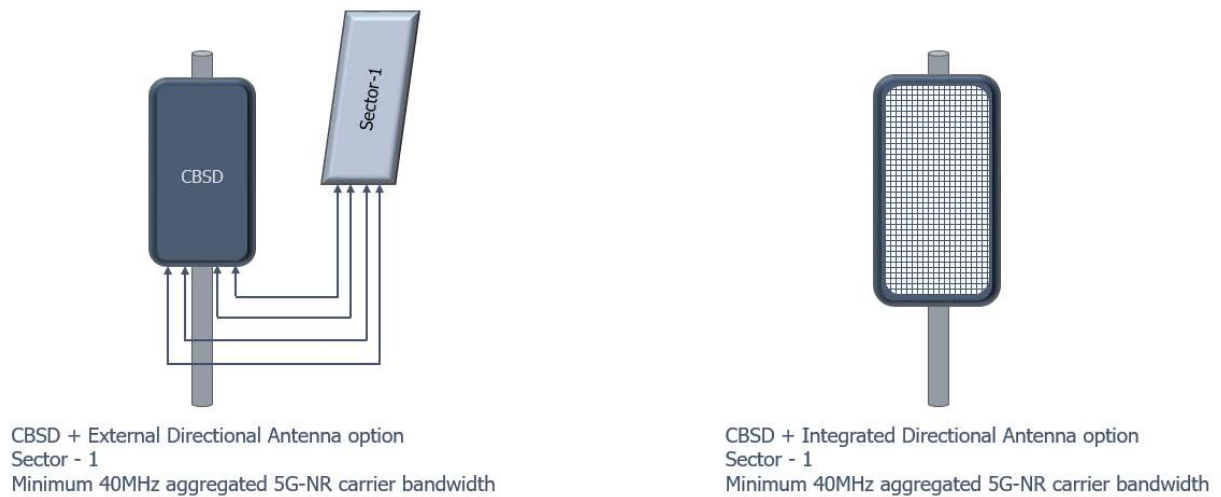
**Figure 8 – Tri-Star Configuration Product**

“Inside-out” could provide outdoor / pedestrian street coverage utilizing its small and medium business customers’ locations. Charter has conducted tests in various SMB locations in New York and Los Angeles markets. Results show that the indoor low power Category-A CBSDs can provide blanket coverage when deployed every 400 – 500 ft.

A strand or attach mount usually compliments SMB coverage by serving as an umbrella cell and fill in any coverage holes. A Tri-Star configuration which adds a third sector to cover indoor and two sectors pointing outside the stores from behind a glass window is being evaluated. Operators can make use of this deployment type where applicable to form a uniform layer of CBRS coverage targeting an areas of interest (AOI) and have run trials confirming seamless connectivity and performance between them.

#### **4.1.3.4. Attached Mount**

In high demand markets such as New York, operators may leverage locations where it has wireless rights on the buildings to deploy high power Category-B CBSDs for outdoor mini-macro type deployment. Attached mount nodes can be installed more strategically than strand mount or SMB to clear obstructions or point to a targeted hotspot with required down-tilts like other two types (strand, SMB) of deployment. Attached mount deployment use case brings possibility to deploy most advanced antenna features such as MIMO and antenna beamforming to support 5G technology deployment and provide high end user performance speeds. This type of deployment provides larger coverage and capacity than strand and SMB scenarios.



**Figure 9 – Attached Mount Mini-Macro**

#### **4.1.3.5. Indoor Enterprise**

This scenario covers indoor deployment of CBSDs into large enterprises, multi-dwelling units (MDU) and indoor venues. Charter has done indoor enterprise trials in its Spectrum Plaza office building with 50 indoor Category-A type CBSDs with 4G CBRS service.

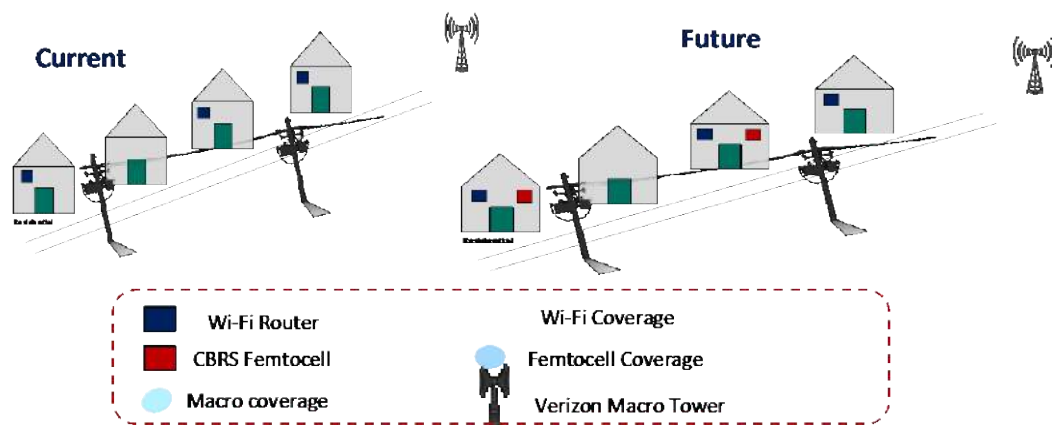


**Figure 10 – Example Cat-A CBSD**

#### **4.1.3.6. Residential Femtocell**

For a typical residence, a femtocell shall provide comprehensive coverage inside and around the house.





**Figure 11 – Residential Femtocell Deployment**

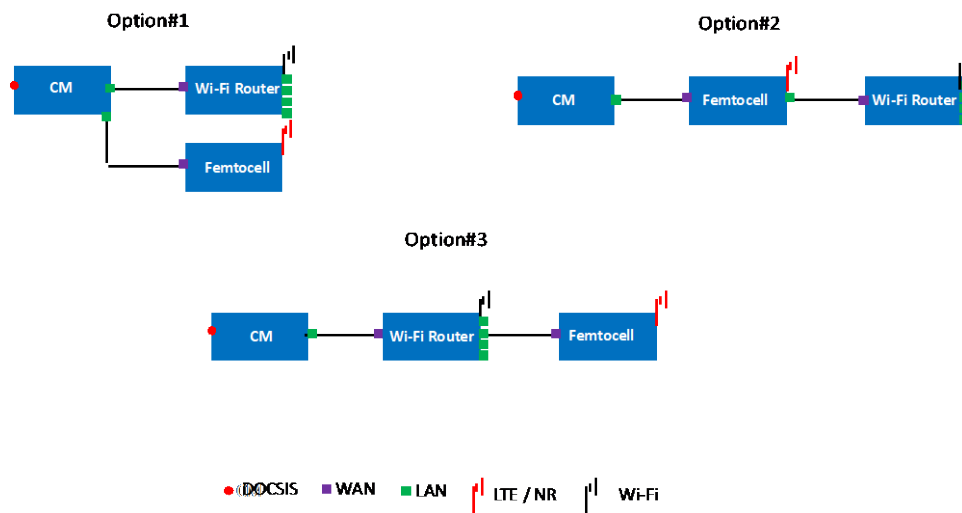
Key Features of Residential Femtocell scenario are as follows:

- Deployment Type: Indoor
- CBSD Class: CAT A
- EIRP: 30 dBm
- Antenna: Integrated Omni
- Backhaul: Ethernet/DOCSIS
- Power: AC

#### **4.1.3.7. Residential Femtocell and In-home Connectivity**

MSOs have a range of options for standalone femtocell connectivity in the home, which are depicted in Figure 12. Each option has its advantages and disadvantages. Option #3 in Figure 12 is intuitive and used by a couple of MNOs in the USA, but can it support TDD LTE timing and synchronization requirements without requiring hardware upgrade on the Wi-Fi router? Option #1 will require the addition of new LAN ports on the CM. Option #2 will put the femtocell in the path of all the broadband traffic in the home, which may not be optimum.

The solution the industry selects for timing and synchronization will be one of the primary factors in deciding the in-home connectivity model. The timing and synchronization topics are covered in Section 5.3 of this paper.



**Figure 12 – Residential Small Cell Gateway Options**

#### **4.1.4. Backhaul**

The near-ubiquitous availability of cable and DOCSIS assets in urban and suburban areas may be one of the key enablers for the industry to deploy small cells at scale economically.

There are, however, some preparation and planning around bandwidth, latency, QoS, and timing that operators need to make.

##### **4.1.4.1. Bandwidth**

Depending on the wireless node capabilities, configuration, and network architecture, the peak and average bandwidth demands on the DOCSIS backhaul will vary.

For example, a 20 MHz TDD LTE small cell will demand much less peak bandwidth than a 40 or 100 MHz 5G NR small cell. Similarly, an integrated small cell architecture will require less average bandwidth per node for control plane traffic than a vRAN based architecture.

##### **4.1.4.2. Latency**

Latency is a critical factor in determining the quality of experience for the end-users. High end-to-end latency for user plane traffic can deteriorate user experience enough to render some applications unusable. Additionally, high latency for the control plane traffic may break important cellular features such as seamless handover.

MSOs are evaluating the application of CableLabs' Low Latency Xhaul (LLX) [6] technology in reducing latency through mobile and DOCSIS scheduler pipelining.

#### 4.1.4.3. Quality of Service (QoS)

The HFC plant is a shared medium used by many customers and applications simultaneously in the time and frequency domain. Luckily, DOCSIS offers an extensive set of tools (e.g., service flows and classifier) to logically separate and treat traffic differently for different applications.

Layer 2 and layer 3 traffic tagging capabilities to allow the cable modem to apply operator provisioned classifiers and ultimately customized QoS for all traffic originating from and terminating at the small cell.

#### 4.1.4.4. Timing and Synchronization

Unlike Wi-Fi, 4G LTE and 5G NR require stringent synchronization (frequency and phase) of wireless transmissions to avoid interference between uplink and downlink. Since the CBRS spectrum is a shared band, the clock synchronization across base stations of both the same and different operators is critical for full realization of the spectrum and to avoid unwanted interference.

As laid out in Table 4, the timing and synchronization requirements for TDD LTE and 5G NR are especially stringent.

**Table 4 – Timing and Sync Requirements of LTE and 5G**

	Frequency	Phase
<b>4G LTE TDD</b>	$\pm 50$ ppb	$\pm 1.5$ $\mu$ s
<b>5G TDD</b>	$\pm 50$ ppb	$\pm 1.5$ $\mu$ s

These synchronization requirements are documented in 3GPP specifications TS 36.133, TS 36.922, and TS 38.104.

The acquisition of accurate phase and frequency for outdoor small cell deployment is rather straightforward using GPS.

On the other hand, the acquisition of accurate phase and frequency for indoor wireless deployment is much more complex and requires evaluation of multiple options.

As highlighted in the table below, there several options for timing and synchronization. For outdoor deployments, GPS is the most widely used timing source. However, for indoor applications such as femtocell, the combination of precision time protocol (PTP) and DOCSIS Time Protocol (DTP) ranks higher on the list and is being carefully studied and tested.

**Table 5 – Options to Provide Timing and Sync**

	Advantages	Disadvantages
<b>DOCSIS Time Protocol (DTP) with PTP</b>	<p>Supports LTE TDD and 5G timing precision requirements</p> <p>Timing from operator-owned and operated network</p> <p>CableLabs standard promoted by cable vendors</p>	<p>Requires significant changes to DOCSIS infrastructure, including hardware upgrade to CM</p> <p>Grand Master clocks in each headend</p> <p>Regular network calibrations MAY be required</p>
<b>Global Positioning System (GPS) *With and without network assist</b>	<p>No upgrades to DOCSIS network required</p> <p>Supports LTE TDD timing precision requirements</p>	<p>Receive challenges indoors, susceptible to jamming</p> <p>Placement not in the control of the operator</p> <p>Installation and operation cost external antennas</p>
<b>Network Listen/Macro Sniffing (e.g., synchronization signals from Macro cells)</b>	<p>No upgrades to DOCSIS network required</p>	<p>Reliance on macro network for timing, availability everywhere could be an issue</p> <p>Out-of-band listen requires dedicated radio – additional cost &amp; more space</p>
<b>PTP over-the-top</b>	<p>No upgrades to DOCSIS network required</p>	<p>Timing synchronization not precise enough for TDD LTE even with DOCSIS QoS. (5-10 millisecond range)</p> <p>Performance is negatively impacted with network loading and uplink packet delay variation (uplink bandwidth limited)</p>
<b>Network Time Protocol (NTP) over-the-top</b>	<p>No upgrades to DOCSIS network required</p>	<p>Timing synchronization not precise enough (100 millisecond) even with dedicated QoS on DOCSIS</p>
<b>TV Broadcast Listen</b>	<p>No upgrades to DOCSIS network required</p>	<p>Need a receiver for TV broadcast</p> <p>Femtocell must know its own location &amp; TV tower</p>

#### 4.1.5. Dual-SIM

##### 4.1.5.1. What is DSDS

Dual SIM dual standby or DSDS is a device that has 2 SIMs installed into it. Both SIMs are active only when not in a call but when the user places a call, the other SIM is placed in standby. In the case where a service provider (SP) wants to use one SIM subscription for voice and text and the other SIM subscription for data, updates to the device firmware are required for the reconfiguration.

##### 4.1.5.2. Charter Use Case

MSOs are typically in a MVNO relationship with an MNO. Charter wants to improve the MVNO economics. Since Charter is not providing voice or text on its RAN networks, at this time, Charter requires the device to scan and connect to both networks at the same time or switch back and forth as needed. This requires two subscriptions which entails the need for two SIMs.

##### 4.1.5.3. How Does DSDS Work

DSDS is a derivate hybrid between dual SIM single standby (DSS) and dual SIM dual active (DSDA). In DSDS, both SIMs are in standby mode as long as neither is in a call nor actively listens for paging messages from both networks while idle. However, once a call is received on one of them, the other SIM becomes inactive and is unable to receive calls or messaging. When incoming calls come in for the inactive SIM they are simply routed to voicemail and SMS messages are held until the active SIM goes idle.

**Table 6 – Comparison of Three Dual-SIM Technologies**

	Dual SIM Dual Standby	Dual SIM Single Standby (Passive)	Dual SIM Dual Active
Overview	A hybrid between Dual SIM Standby and Dual SIM Dual Active	Worst implementation of Dual SIM technology, and affordable phones commonly use it	It can also receive calls on either of the two SIM cards, at the same time
Technology	Smartphone have two active SIM cards, and they both use only one radio transceiver. However, they are both active only as long as they are not used, hence the name of Dual Standby.  As long as the SIM cards are both in Standby mode, calls can be made and received on any of them. However, once a call on one SIM card is taken, the other becomes inactive and the first card is no longer actively used.	Capable of using two different SIM cards, but only one of them can be active at any time, hence when one SIM card works, the other is unreachable.  To use the second SIM card, it needs to be manually activated and the first SIM deactivates	Both the SIM cards are permanently active. During a conversation on one of them, the other still works and receives calls, messages or data.  The disadvantage of Dual SIM Dual Active phones is that the devices have two radio transceivers, one for each SIM card, hence consuming more battery than regular Single SIM smartphones

Usage	Most popular implementation	Used in older mobile phones (not smartphones) and lower priced phones	High battery consumption and more expensive to manufacture, leading to a higher device price
-------	-----------------------------	-----------------------------------------------------------------------	----------------------------------------------------------------------------------------------


#### 4.1.5.4. How Switching Decisions Are Made

In normal DSDS, the device will not switch on an active data session until it is no longer able to maintain connection to the network. Charter has worked with the original equipment manufacturers (OEMs) to switch based on several factors. Signal strength and signal quality are used to determine the need for the device to switch networks. When a device is out of coverage of operator's CBRS network, it maintains data, voice and text on the MNO's network. This allows the user to maintain their experience at all times.

When a device enters a geographic local, geo-fenced area, the device enables the second SIM and the device begins to actively search for a suitable network. Once the user equipment (UE) determines that the level, quality and hysteresis timer has elapsed, it will then preform an attach request to the network. The UE will maintain normal mobile operations with the one exception that it continues to listen to the first network for incoming pages for text and voice calls.

When the device receives a voice call from the MNO network, it then answers this request and reattaches to the data network of the MNO so the user may continue to use data. This is done due to the fact that the UE can still only have one active network on a DSDS device. When the call ends, the device again searches for the MSO network and, if available and criteria met, reattaches the data stream.

As the device begins to leave the MSO coverage area, signal quality and strength reduce. Rather than allowing it to lapse into a radio link failure (RLF), the device moves back to the MNO network if available, based on predetermined exit levels of strength and quality.



**Figure 13 – How a DSDS Device Moves Across Networks**

#### 4.1.5.5. Dual eSIM

Embedded-SIM or eSIM is an embedded universal integrated circuit card that allows a user to store a cellular profile without the need for a physical chip that needs to be inserted into the device. It allows for the storage of multiple SIM profiles to be stored, but at this time only one can be used at a time and the user

has to manually switch between each profile to use it. New features will need to be developed to allow two eSIM's to function as the current physical SIM (pSIM) and eSIM.

#### 4.1.5.6. *Is DSDS Good Enough*


With the modifications to the operating systems and the chipset enhancements that have occurred over the last year, DSDS functionality is very good. As it stands to date, the user would generally not realize as the device moves between the networks. Working with the chipset vendors and the device manufacturers, there may be some room for improvement. Regardless, the technology today will satisfy many needs for private network owners who wish to connect or offload data to their network. This allows them to control the access to their network with increase security and still allow calls and messages to be received without having to set up their own voice and data network.

#### 4.1.6. *Virtualized RAN Architecture*

Charter's medium to long-term network vision is to transition from standalone RAN to a vRAN architecture where baseband functionality is centralized and virtualized, and supports standard open interfaces. Ultimately, Charter is driving for flexible hardware and software implementations affording scalable, cost-effective network deployments.


The transition to a vRAN based architecture is being evaluated for all the mobile-offload deployment scenarios described earlier in the paper. For emphasis, the use of vRAN is not only being evaluated for outdoor deployments but also for indoor residential deployment using femtocells.

For the near-term deployments, an operator's focus can be on the integrated solution, which includes both the Radio Unit (RU) and baseband unit (BBU) functionality in a box at the edge.




**Figure 14 – vRAN Architecture – Backhaul**

The long-term RAN architecture aims to centralize and virtualize the Centralized Units (CU) functionality, leaving the RU, and Distributed Units (DU) features at the edge of the RAN. Operators also want the flexibility to dynamically allocate DU functionality either at the edge of the RAN or centrally depending upon the backhaul option available.



**Figure 15 – vRAN Architecture – Midhaul**

When fiber backhaul is available, operators can consider two options: collocating the DU/RU to align with the DOCSIS deployment scenarios or to move the DU to a centralized location.



**Figure 16 – vRAN Architecture – Fronthaul**

The mobile community envisions improved RAN performance with vRAN architecture compared to a traditional, integrated RAN. For example, the performance enhancements would be realized by utilizing better interference management and improved mobility. In the long run, there are opportunities to reduce total cost of ownership (TCO) from baseband pooling.

#### **4.1.7. Fixed-Mobile Convergence**

Charter has years of experience designing, deploying, and managing Wi-Fi networks both outdoors and indoors, leveraging Charter’s extensive DOCSIS and PON networks. Charter will expand its wireless networks with the deployment of 5G NR small cells. Charter’s availability of DOCSIS and PON assets in urban and suburban areas is one of the key enablers allowing Charter to deploy 5G NR small cells at scale economically.

However, fixed mobile convergence is more than network convergence, it is about customer experience. Fixed mobile convergence allows Charter to provide a single cohesive experience to its customers. Convergence at the service level consolidates subscriber management and policy enforcement allowing customers to enjoy their services as they move between access networks. A parent can assign parental controls to a child’s device and those controls work whether the device is on the cable broadband network or on the mobile network. A premium subscriber would get the same service at the neighbor’s house or the local coffee shop, even if those locations have lower service. With a converged service, the network would be able to identify and enforce network policies on customer devices no matter where they are and give them the same service and experience.

Charter’s vision for fixed mobile convergence is to deliver ubiquitous wired-wireless connectivity to our customers anywhere and on any device. Customers would carry their services, policies, and identity with them wherever they go. Through Charter’s unique broadband and wireless assets, this can be delivered through the deployment of high capacity and low latency networks.

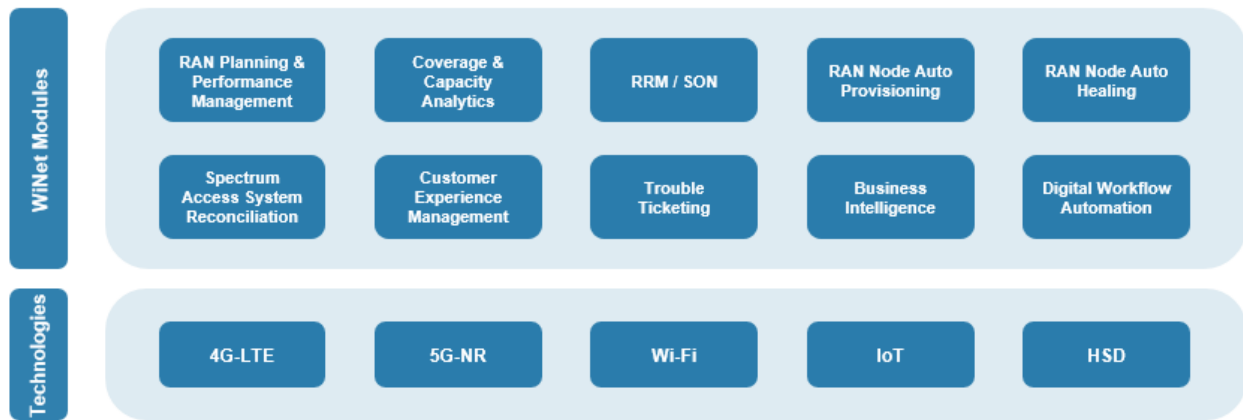
#### **4.1.8. Automation**

Charter envisions a new automation tool in its production network for its commercial network for applications such as network automation intelligence and workflow orchestration, network integration and network management. The tool enables end-to-end automated network life-cycle management for a large-scale network rollout and also enables automated reporting and provides a high level of customization. The tool’s customizable dashboards and dynamic widgets provide end-to-end project visibility for seamless status tracking of network elements, work orders and tasks. It’s a real-time, technology-agnostic single platform for live network analytics by leveraging data points from multiple sources for an enablement of valuable insight about network’s health.

The tool is closely integrated with different network elements for an end-to-end network visualization, configuration management, performance management, fault management, change management features. It enables not only proactive network diagnostics but also provide remediation for performance enhancement. It is scalable platform with fully developed infrastructure to handle massive data volumes and transactions,



allowing an operator to handle continuous growth in user data due to high bandwidth speeds and an increase in number of mobile devices.



**Figure 17 – Charter Wireless Automation Framework**

## **4.2. Cox Communications**

Convergence with wireless is currently centered around leveraging our fiber and HFC plant to provide services to the MNO's that include: Fiber To The Macro Cell Site (FTTX), strand mount small cells on our aerial HFC plant that incorporated AC power and DOCSIS backhaul transport, and offering CRAN (Centralized Radio Access Network) facilities space and fiber transport to upgrade the MNO's architecture as they move toward cloud native RAN.

Cox believes the market is becoming more attractive for us to enter the wireless space and we are exploring it more aggressively now but have not announced any specific plans.

As Cox looks forward to evaluating the current wireless segment, we look back to take the lessons learned from our earlier wireless attempts of both an MVNO and MNO solution. The challenge 10 years ago was in understanding the customer demographics with the then recently launched iconic smartphone device that limited our differentiation, as well as having wholesale metered voice, text and data on a much slower and costlier 3G network from our MVNO provider. With the trial markets we built in our cable franchised footprint, we wanted to provide ubiquitous service and our subscriber penetration rates couldn't support the cost of the wireless infrastructure as well as having to work roaming agreements without the scale of a large MNO. We are more optimistic with today's wireless environment as other cable operators have successfully launched MVNO's using high speed LTE networks and having driven their wireless offload tonnage by Wi-Fi both in the home and metro systems as well as evaluating other solutions such as small cells to drive the wholesale usage costs down.

As we look to the future, there are several fronts we see opportunity that include: driving Low Latency Xhaul (LLX) that enables low latency 5G wireless backhaul, looking at fixed wireless access that extends the edge of our HFC and fiber plant to add incremental households as well as enabling the opportunity to support rural broadband deployments, deploying Wi-Fi 6E in our wireless gateways that we can drive higher data throughput experience to our customers over our broadband network, and finally as we explore opportunities with retail wireless.

### 4.3. Shaw Communications

Shaw Communications is a Canadian cable and mobile operator that provides both wireline services to over four million homes and businesses in western Canada, as well as mobile services to customers in British Columbia, Alberta and Ontario. Shaw has offered broadband services in Canada since 1996 and became a mobile operator in 2015 with the purchase of Wind Mobile, which it has since rebranded to Freedom Mobile. In July 2020, Shaw launched the Shaw Mobile brand to provide a new wireless service that leverages Shaw's Fast LTE and HFC/DOCSIS networks to provide Shaw Internet customers with an innovative wireless experience.

Over the past five years, Shaw has invested heavily in its wireless network, adding macro sites and purchasing new spectrum in order to improve the performance and reach of its rapidly growing wireless service. Since its entry into wireless, Shaw has focused significant effort on leveraging synergies between its established wireline business, and its new and growing wireless division. Indeed, access to Canada's largest public Wi-Fi network (Shaw Go WiFi), fiber backhaul, critical facilities, buildings, operational teams, retail, as well as many other opportunities proved to be of significant value early on. However, the ongoing exponential growth in wireless traffic and the impending arrival of 5G has put network and service convergence as a key focus area.

Seamless connectivity is the foundation for our future economy, but this future will require extremely close interplay between mobile and fixed technologies. The deployment of the 5G vision will require an unprecedented level of network connectivity and densification, as well as previously unseen levels of collaboration between wireline and wireless technologies [7].

In order to deliver both 5G as well as 10G (multi-gig fixed broadband), converged Canadian operators like Shaw will need to find ways to leverage their key strategic advantages to compete with their well-funded and converged Canadian Telco competitors. The top three strategic opportunities are:

- **Hybrid-Fiber Coax Infrastructure** – Already deployed to virtually every house and business in the country, HFC infrastructure makes an ideal solution for the densification of wireless networks. Able to transport multi-gigabit traffic today, it can also transport power, greatly reducing the time and cost of small cell deployments.
- **Hub Site Facilities** – Cable's historic hub-site topology, gives operators access to significant spare power and cooling, just a few miles from the customer. Thanks to the new "Distributed Access Architecture", these new facilities will soon be vacated, giving cable operators access to a network of distributed mini-datacenters, which can then be leveraged to provide ultra-low latency virtual/cloud RAN services, which are likely to be a key component of 5G roll outs.
- **Core/Service Convergence** – Both wireless and wireline networks leverage virtually the same architecture, leveraging a "wireless core" or "CMTS" to control the flow of wireless or wireline traffic. The convergence of these two cores will ultimately enable a truly seamless and differentiated experience for customers, increasing security, flexibility and control.

The opportunities described above are indeed highly strategic to virtually all cable operators globally. However, in order for them to be realized, several key challenges must be addressed as an industry. These challenges include:

- **Latency Reduction** – 5G as well as new mid/fronthaul solutions will require ultra-low latency, which cannot be easily met by today's consumer focused shared access networks like DOCSIS and PON. Improvements in new DOCSIS technologies, such as Low Latency Xhaul (LLX) [6], will be required

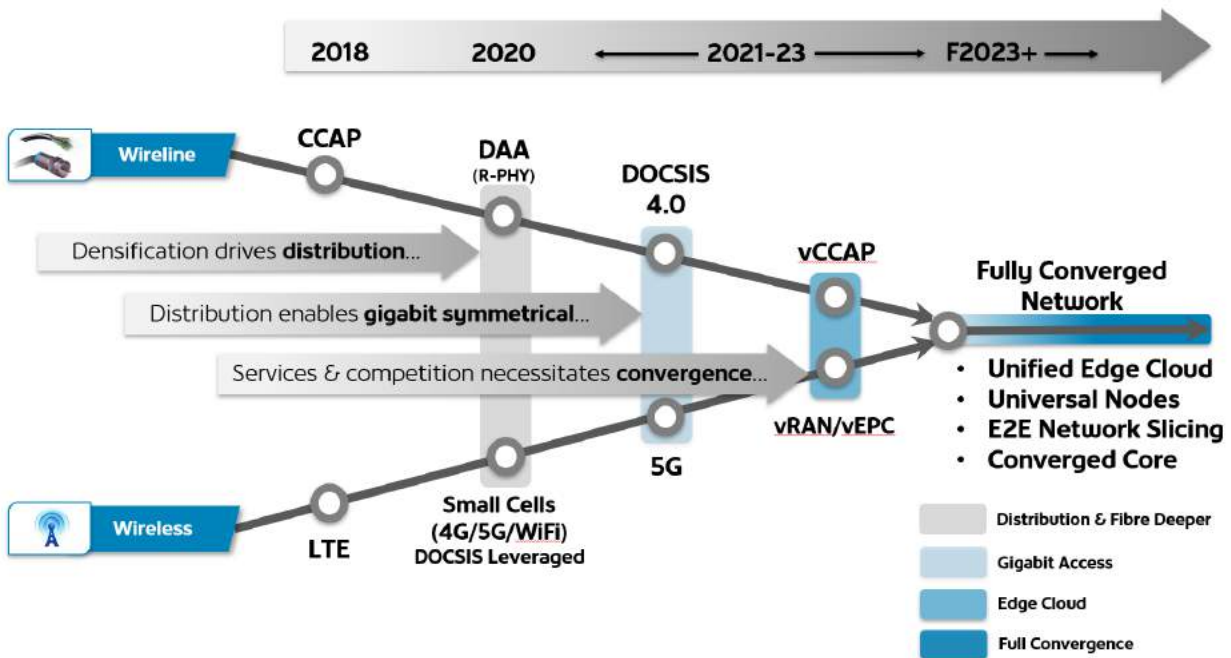
in order to reduce latency to sub 1ms level, making HFC virtually indistinguishable from dedicated fiber for 5G and mid/fronthaul purposes.

- **Timing Distribution** – Delivery of highly accurate timing to small cells is challenging using today’s DOCSIS protocol. However, innovations such as the DOCSIS Time Protocol (DTP) [19] will solve this challenge, eliminating the need for costly GPS integration, and greatly improving the flexibility and efficiency of small cell installs.
- **Xhaul Over HFC** – Today’s fronthaul protocols are extremely demanding and designed for dedicated fiber use. In order to enable the ultra-dense and demanding wireless networks of the future, new “xhaul” techniques must be developed that will allow operators to leverage the efficiencies of fronthaul while using shared access topologies such as HFC and PON.
- **Core Convergence** – While similar, wireless and wireline core technologies differ in numerous ways. Converging these technologies to enable a host of benefits for consumers will require development and integration across a diverse vendor ecosystem.

While challenging, the issues outlined above are certainly possible to overcome, and the industry is already making progress to address these. Recent models built at Shaw, indicate that a cable operator could reduce its build cost by up to 99% and build time by up to 95% by leveraging existing HFC infrastructure, rather than extending fiber to feed small cell densification [8]. The promises of cloud RAN and core convergence also carry similar promises of greatly reduced costs and improved customer services. However, operators will need to rally behind these new technologies and drive their development in order to see the benefits realized.

Ultimately, we believe that both cable and wireless technologies are evolving in lock step, and our industry is exceptionally well positioned to realize the incredible synergies between two of our most important assets. As shown in Figure 18 below, today’s exponential growth in broadband is necessitating densification of our networks, accomplished through distributed access architecture (DAA) in wireline and small cell deployment (ideally enabled through DOCSIS) in wireless. That densification is also an enabler for multi-gigabit symmetrical services, enabled through DOCSIS 4.0 in wireline, and 5G in wireless.

In the future, competitive and efficiency drivers will require operators to embrace cloud and virtual network infrastructure, enabled by virtualized CCAP (vCCAP) in wireline and vRAN/vEPC in wireless. Shaw believes all of these steps will happen in very close timing, which will lead to the ultimate convergence of both network infrastructures. The final converged state will enable a myriad of new opportunities and efficiencies, such as the unified edge cloud, universal nodes, end-to-end network slicing and orchestration, converged wireless/wireline cores, new services and much more.



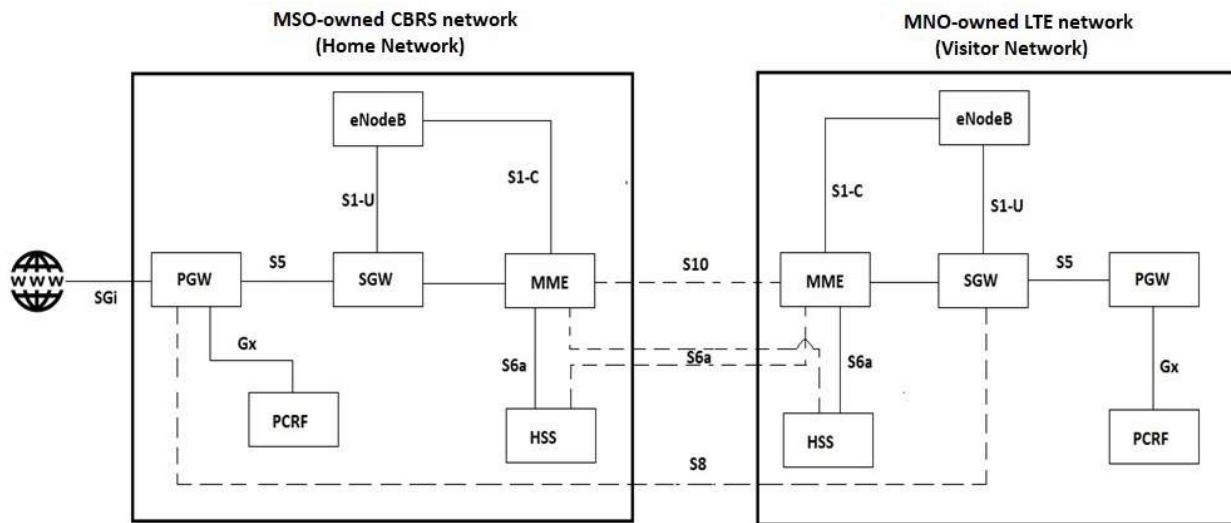
**Figure 18 – Wireline and Wireless Convergence Roadmap**

In addition to new technology changes, Shaw strongly believes that several other industry challenges must be overcome in order to ensure Canadians have access to world leading wireless services and 5G. These include inter-operator handover, access network convergence, and Wi-Fi interoperability, and will be discussed in more detail in the sections below.

#### **4.3.1. Inter-Operator Handover**

Existing regional MNOs, like Shaw, typically rely on domestic roaming agreements to provide coverage outside of their home network footprint. With traditional roaming interfaces, however, subscribers often experience network problems such as dropped calls or interrupted downloads/uploads as they move between the home network and the roaming partner's visited network. This negatively impacts customers' perceptions of network quality and brand, which undermine a regional operator's ability to effectively compete and grow in a new market. This will also be the case for MSOs planning to deploy CBRS, where the initial coverage areas will typically be more localized.

Fortunately, the 3GPP Home Routing (HR) specification provides a standards-based approach for supporting seamless inter-operator handover between the operator's home network and its roaming/MVNO partner's visited network. With HR, the data plane traffic is routed back to the home network, giving the home network operator control over the subscriber's traffic and its IP context. This allows subscribers to roam between the home and visited networks without service interruptions. Figure 19 shows the network architecture for HR.



**Figure 19 – Home Routed (HR) Network Architecture**

To implement the HR model, the home and visited networks are required to share three interfaces: S6a, S8, and S10. These are described briefly below.

- S6a—is an interface between the Mobility Management Entity (MME) and the Home Subscriber Server (HSS) of both networks that enables the transfer of subscriptions for authenticating and authorizing user access to the network.
- S8—is an interface between the Serving Gateway (SGW) of the visited network and the Packet Gateway (PGW) of the home network, acting as an inter-Public Land Mobile Network (PLMN) reference point to transfer user traffic back to the home network. S8 allows the home network to control a subscriber's traffic even when the subscriber is roaming on the visited network.
- S10—is an interface between two MMEs used for bearer modification with MME relocation and MME-to-MME information transfer. S10 enables seamless data session transfer in connected mode. In addition to sharing roaming interfaces, HR implementation requires each network to be configured with mobility parameters that utilize connected- and idle-mode triggers.

The HR implementation is ideal for operators that have a roaming/MVNO agreement with an MNO that allows sharing of the roaming interfaces (i.e., S6a, S8, and S10) and control over mobility configuration. In addition to establishing the roaming interfaces, both operators also need to exchange cell site information to update their neighbor relations tables that identify the other operator's adjacent cell sites to which handovers occur.

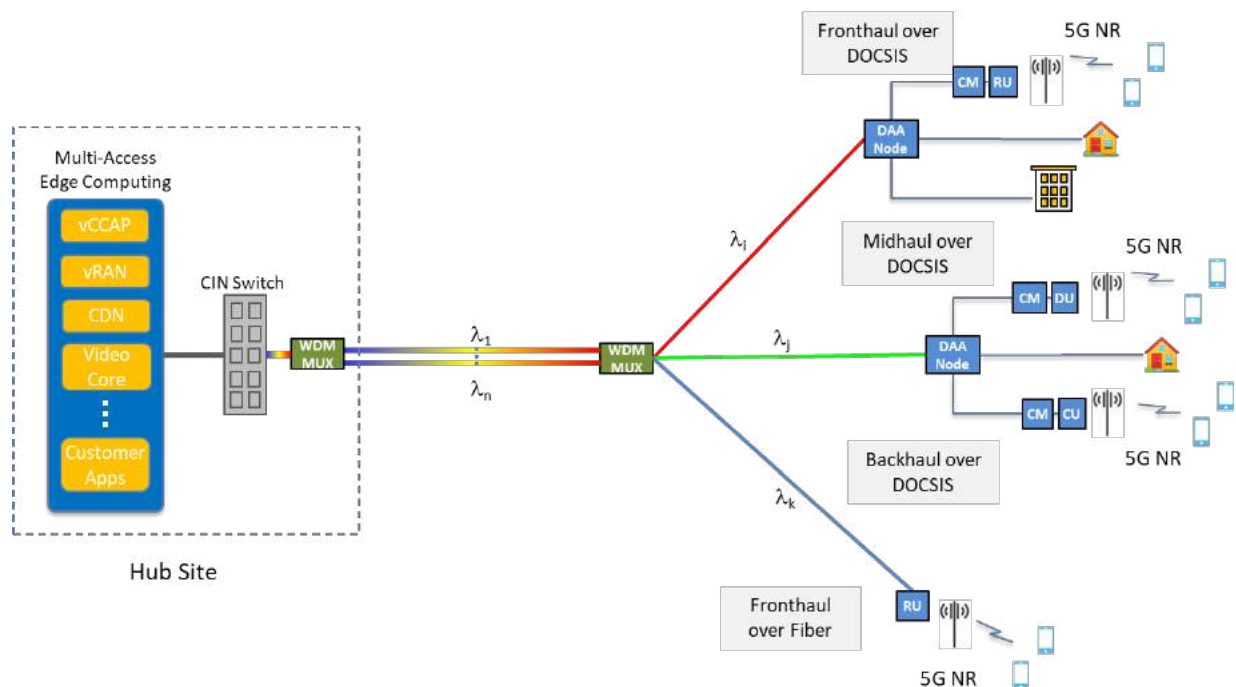
Inter-operator mobility allows regional operators/MSOs to provide high quality services to their subscribers on par with the national wireless carriers as they expand their networks into new markets. Adopting inter-operator mobility would also promote competition, investment, and network deployments by regional operators, particularly in rural and remote areas. With inter-operator mobility, regional operators will have the ability (and incentive) to expand into smaller communities and rural areas more quickly and compete with the incumbents more effectively because they can ensure high quality service to their customers as they build out their network footprint.

### 4.3.2. Access Network Convergence

5G deployments will drive a dramatic increase in network density through small cell deployments. Although dark fiber has been used extensively in the past to support 3G/4G macro site deployments, unprecedented levels of new fiber builds would be required to support the roll-out of 5G. Fortunately, MSOs have already been deployed HFC/DOCSIS access networks down virtually every street and to almost every building on the continent to the point where today it reaches 93% of American households. And with the recent release of the DOCSIS 4.0 standard, HFC networks will soon be able to deliver multi-gigabit capacity.

As such, MSOs are uniquely positioned to create ultra-dense 5G small cell deployments by leveraging their existing hybrid fiber coax (HFC) networks and emerging technologies such as DOCSIS 4.0, DAA, LLX, and multi-access edge computing. The main advantages of HFC/DOCSIS networks relative to other alternative networking technologies (e.g., dark fiber) are its low cost, scalability, access to power, and ease of deployment. For the past couple of years, Shaw has utilized its aerial HFC plant to deploy LTE small cells for additional capacity and/or coverage in targeted areas. Deploying small cells on aerial plant addresses three major challenges with outdoor small cell deployments: site access, backhaul and power. That is, site access is usually already covered by existing pole-line attachment agreements and both backhaul and power are provided over the coaxial cable plant.

Figure 20 shows as example of a converged access network architecture with several options for transporting 5G small cell traffic over the HFC/DOCSIS network.



**Figure 20 – Converged Access Network Architecture**

As shown in the figure above, 5G fronthaul, midhaul and backhaul can be viably transported over DOCSIS using the technologies mentioned earlier (e.g., LLX, DTP). With fronthaul, the 5G remote unit (RU), which implements basic RF functions (e.g., filtering, amplification), is located at the cell site and the other radio functional blocks are implemented within the virtualized RAN (vRAN) at the hub site. Fronthaul offers the

highest performance and efficiency but has the most stringent latency and throughput requirements. Midhaul and backhaul implement different functional splits, at what are known respectively as the distributed unit (DU) and centralized unit (CU). These splits have less demanding transport requirements but provide lower performance and efficiency compared with fronthaul.

In cases where the small cell is already near existing fiber plant, fronthaul can also be carried over fiber by assigning a dedicated wavelength to the small cell and sharing the common wavelength division multiplexing (WDM) link between the hub site and a remote WDM multiplexer (MUX). The preferred transport option depends on several factors such as the latency requirements of the end user applications, the available 5G spectrum (and hence capacity) the radio site, and the cost of deployment.

The other opportunity for access network convergence is at the hub site (or headend). Both the virtualized CCAP (vCCAP) and vRAN can be implemented on multi-access edge computing platforms in which compute and storage resources are shared across multiple applications. CableLabs and other standards bodies are developing architectures and specifications to realize this important opportunity.

With the fundamental technologies (e.g., DOCSIS 4.0, DAA, LLX, DTP, multi-access edge computing) needed for wireless/wireline access network convergence still at relatively early stage of development, broad support is needed across the industry to make them a commercial reality.

#### **4.3.3. *Wi-Fi Interoperability***

Although 5G mobile networks will be deployed extensively throughout the world, Wi-Fi will continue to carry the bulk of Internet traffic well into the future. According to Cisco's latest Annual Internet Report [9], public Wi-Fi hotspots are expected to grow four-fold from 2018 to 2023. As such, public Wi-Fi networks represent an important opportunity for MSOs to offload mobile data traffic from their own mobile networks or those of their MVNO partners to reduce network build and/or roaming costs.

Over the past decade, Shaw has built Canada's most extensive service provider Wi-Fi network, Shaw Go WiFi, with over 100,000 public access points deployed to date. Shaw Go WiFi extends our Internet customer's broadband experience beyond the home as a value-add to our customer's wireline network experience. Over 3.6 million devices have authenticated on the Shaw Go WiFi network, which is used by our customers in coffee shops, restaurants, gyms, malls, public transit and other public spaces from British Columbia to Ontario.

Freedom Mobile and, more recently, Shaw Mobile customers are also able access to the Shaw Go WiFi network along with over 300,000 home hotspots deployed in our Internet subscribers' homes across the country. Our mobile customers' devices automatically connect and authenticate to our public and home hotspots via Hotspot 2.0 (aka Wi-Fi Certified Passpoint) using their wireless credentials stored on the SIM. This provides mobile customers with extended network coverage and offloads a significant proportion of traffic from the mobile network.

With the introduction of Wi-Fi 6 (IEEE 802.11ax), dramatic improvements are expected in Wi-Fi capacity, efficiency, and reliability in the coming years. The Cisco Annual Internet Report predicts that Wi-Fi 6 hotspots will grow 13-fold from 2020 to 2023 and represent 11% of all public Wi-Fi hotspots by 2023. The recent allocation of 1.2 GHz of unlicensed spectrum in the 6 GHz band in the U.S. also has huge potential for Wi-Fi and will enable the Wi-Fi 6 evolution by alleviating congestion in the existing 2.4 GHz and 5 GHz unlicensed bands. The first wave of unlicensed devices capable of leveraging 6 GHz is expected in the U.S. in the final quarter of 2020 and 60% of devices are anticipated to be Wi-Fi 6 capable by 2022.



3GPP Release 16 also promises to further enhance mobile/Wi-Fi network convergence through Access Traffic Steering, Switching and Splitting (ATSSS). Based on Multipath TCP (MPTCP), ATSSS allows separate 5G NR and Wi-Fi traffic flows to be simultaneously established between an ATSSS-capable mobile device and the core network, providing highly available and robust services. At present, ATSSS is an optional feature for user devices and the 5G core network so operator support is vital to its success.

## **4.4. Vidéotron**

### **4.4.1. Who is Vidéotron?**

Vidéotron ([www.videotron.com](http://www.videotron.com)), a wholly owned subsidiary of Quebecor Media Inc., is an integrated communications company engaged in cable television, entertainment, Internet access, cable telephone and mobile telephone services. Vidéotron is a leader in new technologies with its Helix home entertainment and management platform. As of June 30, 2020, Vidéotron was serving 1,497,300 cable television customers, and 472,200 subscribers to its Club illico video streaming service.

Vidéotron is also the Québec leader in high-speed Internet access with 1,753,100 subscribers to its cable service as of June 30, 2020. As of the same date, Vidéotron had 1,404,900 subscriber connections to its mobile telephone service and was providing cable telephone service to 979,600 Québec households and organizations. Vidéotron has been recognized as one of Montréal's popular employers.

In 2010, Vidéotron built and launched its own 3G (AWS) network and upgraded it to LTE in 2014.

In 2018, Vidéotron Launched Fizz Mobile, 100% digital mobile service brand<sup>1</sup>. This virtualized service provider model is based on Digital Business Platform utilizing TM Forum's Open Digital Framework, which includes business process and information frameworks and Open APIs.

Vidéotron is currently preparing the upgrade of the wireless network to serve Vidéotron and Fizz customers with LTE-Advanced and 5G evolutions.

### **4.4.2. 5G Opens a Very Wide Range of Applications**

5G opens the way to many new opportunities. One of the challenges for operators is to select the right markets to address first with this technology.

There are so many applications that can benefit from 5G, operator will have to expand and consolidate some of their "natural skills". For example, those related to entertainment and content distribution that may evolve to massive contents distribution VR/AR or telepresence, to provide a full set of immersive experiences. An example of the 5G set of services is shown in Figure 21.

IoT is an unmissable opportunity but it has such a wide field of applications that service providers will have to build a strong strategy to select the profitable and strategic markets and serve them with the right IoT technologies and services.

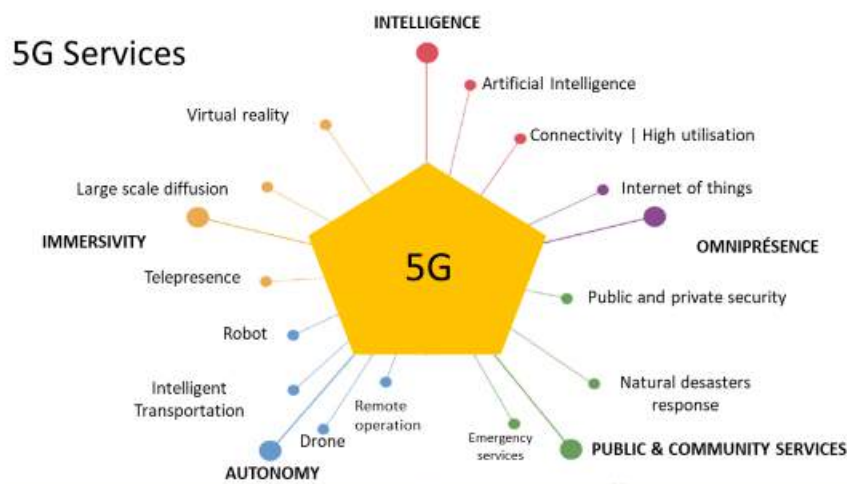
Then many new fields are now being opened to operators with 5G used for example in the "intelligence" field or "autonomy". That creates room for new businesses innovations but also require operators to change their way to do business to be able to propose solutions adapted to this new market in their portfolio.

In this new environment of constantly emerging innovations, operators must consider more than just 5G Radio Access Network and core and include all their network elements in a converged solution to provide a wide range of features to address different markets and services with the right solution.

---

<sup>1</sup> "Vidéotron's disruptive all-digital brand, Fizz, goes from concept to launch in 10 months," by John C. Tanner, Contributing Editor, TM Forum, September, 2019. <https://inform.tmforum.org/casestudy/videotrons-disruptive-all-digital-brand-fizz-goes-from-concept-to-launch-in-10-months/>

Wireless access convergence, fiber network convergence and virtualization are some of the main elements that operators will have to integrate in order to be able to actively participate to this new network panorama.



**Figure 21 – 5G Diversity of Services and Applications**

#### **4.4.3. Wireless Access Convergence – Building on Wi-Fi Assets**

Access to a wireless network coverage is now a fundamental service and all consumers are expecting a good connection at all time for their mobile devices.

But building a mobile network providing a full coverage in any point is a challenge, especially for indoor locations. For LTE networks and for future 5G networks, many solutions exist to improve coverage at any point. For example, small cells to address a specific location that could require better coverage or higher capacity, as well as low band frequencies for long range and better penetration in buildings.

Mobile Network Operators (MNO) are using those tools in the design of their networks, but Wi-Fi is another element that can be very beneficial to wireless coverage, in addition to lowering the cost for MNO and the end user. Indeed, Wi-Fi is now widely deployed in homes, building and many public spaces.

Recent developments in the wireless technologies allow a much greater performance and transparency for users when it is time to decide to connect to traditional wireless network (LTE, or 5G networks normalized by 3GPP) or Wi-Fi Network. Those developments also improve the transition from 3GPP networks to Wi-Fi access point (AP) when the user is moving.

With a strong convergence between 3GPP access and Wi-Fi, it is possible to offer new services or to improve existing services for example:

- Access to a large free Wi-Fi coverage and create a community of users.
- Better customer experience is provided by an automated and seamless connection to the best network available.
- Allow customers to save on their data plans.

- Allow operators to minimize costs as Wi-Fi can be deployed with a lower cost in \$/Bits. This is valid for MNO as they can save on deployment CAPEX & Network OPEX.
- Those technologies also allow agreements between operators for access to Wi-Fi network, including international roaming solutions in order to improve the connectivity for travelers.
- The improved customer satisfaction through the ease of use of Wi-Fi will result in increased customer satisfaction and thus will reduce the level of churn for the operators.

Among all the technologies that enable the convergence between Wi-Fi and 3GPP access network, are:

**Hotspot 2.0 or Wi-Fi Certified Passpoint** (based upon the IEEE 802.11u protocols) now allows the mobile device to connect and authenticate without the need for the user to manually select a network with a very high level of security. This is an essential tool in order to allow the users to connect automatically to the Wi-Fi network (on which they are authorized to access). Hotspot 2.0 thus provides a much easier access to Wi-Fi networks. A single login is required to set the system up and then the system does the rest: searching for accessible networks and gaining entry.

**Access Traffic Steering, Switching and Splitting (ATSSS)** is a feature included in 3GPP Rel-16 to enable high level of convergence between 5G and Wi-Fi networks. This feature that can be hosted in the User Equipment (UE) and the 5G Core (5GC) to provide 3 essential features for convergence:

- Steering: ensure the best network selection between 5G and Wi-Fi
- Switching: allows seamless handover between 5G and Wi-Fi
- Splitting: enables network aggregation to combine capacity of both 5G network and Wi-Fi

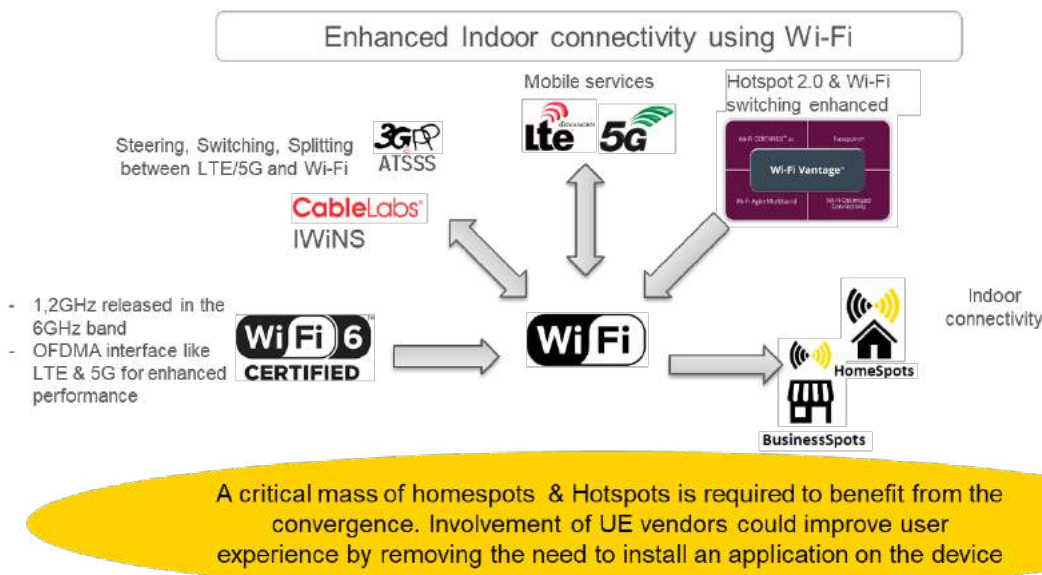
ATSSS allows this advanced management of connection independently for data link and voice service.

**Intelligent Wireless Network Steering (iWiNS<sup>2</sup>)** developed by CableLabs is complementary with ATSSS as it enhances mobile steering and switching between LTE, Wi-Fi and CBRS nodes. iWiNS also has the granularity to apply per-dataflow steering policies by multi-users and multi-networks feedback. This solution requires an application to be installed on the UE and can be improved if the app is integrated into the Operating System (OS) of the device. It can take decisions depending of the state and traffic of the networks (Wi-Fi and LTE/5G) to optimize the communication on a real time basis.

These powerful tools used in conjunction with Wi-Fi 6, the most advanced version of Wi-Fi, will enable wireless access convergence and provide a better connectivity to customers at all times and all places.

---

<sup>2</sup> See “iWiNS Architecture – An aware approach to mobile traffic steering”, Mario Di Dio, Rich Terpstra, CableLabs, August 2019.



**Figure 22 – Wireless Access Convergence**

#### **4.4.4. Network Convergence – All-in-One Fiber Network**

With the arrival of Distributed Access Architecture (DAA), network operators will soon extend their Ethernet network outside of headend and indoor sites to build a new digital network outside plant rather than the current analog amplitude modulation used to transport the digital services (DOCSIS and digital video).

In the next generation of HFC network, operators will deploy DAA and the new nodes will integrate the DOCSIS Physical layer (R-PHY) or the DOCSIS Physical and MAC layers (FMA, Flexible MAC Architecture). In both cases those nodes will be interconnected through an Ethernet digital fiber optics link, sometime called Converged Interconnect Network (CIN).

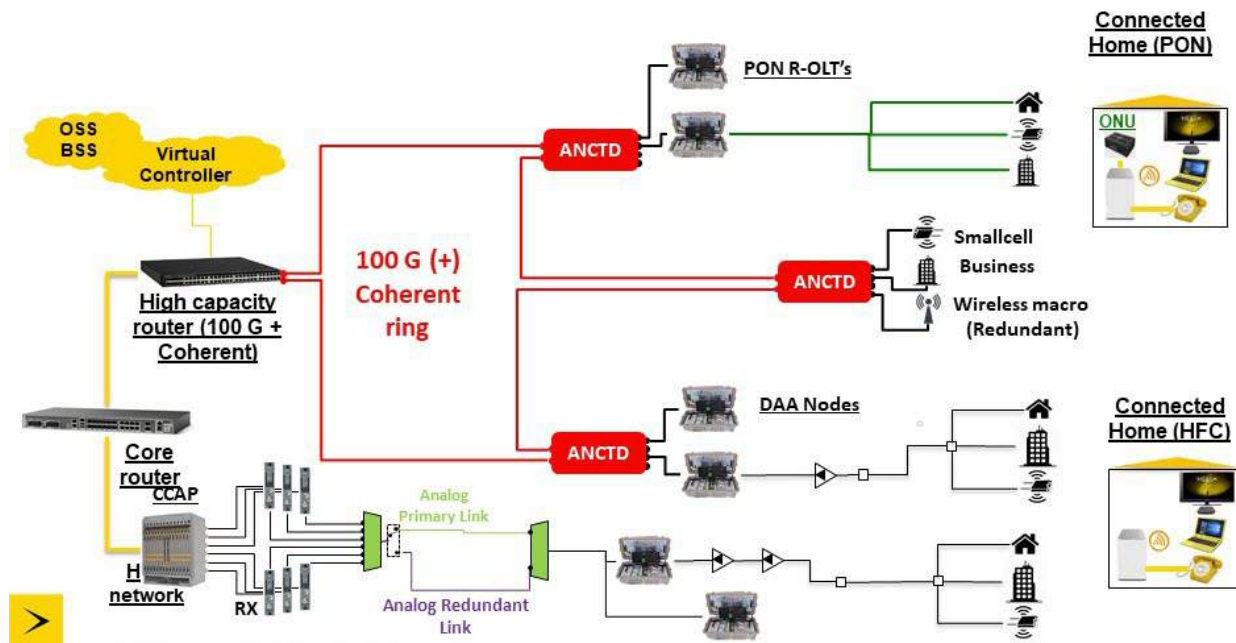
Coax is not the only medium fed by this CIN, some operators will also choose to evolve a part of their network to Fiber To The Home (FTTH) solutions.

Some operators, especially those who already own an HFC network may be attracted by the opportunity to build a PON network (for example 10G EPON) matching the topology of the HFC network. This is now possible using a “remote OLT” (R-OLT) installed in the field and built in a “clamshell” housing like a legacy HFC node or a new DAA node. The position of DAA nodes and PON nodes will then be similar and the topology of digital fiber network (or CIN) used to interconnect DAA nodes and R-OLT will be equivalent.

This CIN being an Ethernet network (typically 10GE), it is possible to aggregate multiple channels (or wavelength) on the same link using DWDM aggregation. In addition to HFC DAA nodes and PON R-OLT, this CIN can also connect some wireless sites (LTE, 5G, Wi-Fi, small cells, macro-sites, hotspots) and some business customers.

As this converged DWDM network will serve a higher number of customers but also a wide variety of services including some with high reliability requirements (for example wireless macro-sites, or business





**Figure 24 – Fiber Network Convergence with High-Capacity Coherent Ring**

#### **4.4.5. Virtualization – A Core Ready for All Opportunities**

Our vision about the virtualization is to create an open ecosystem in which our internal developers and our external partners will be able to integrate their solution. Our foundation uses different sets of technology and hardware based on different standard: ETSI, TM Forum and MEC. We strongly believe that each virtualization technology will benefit Videotron to implement new services and facilitate and support our innovation process.

Hybrid cloud strategy allows deployment of a new business support systems (BSS) system in the public cloud as well as the 5G RAN components in private cloud. The success of the OSS 5G RAN deployment in less than a week has proved that having the right technology in place has improve time to market. Next achievement is to make sure private cloud is up to run network function virtualization (NFV) virtual RAN for the 5G radio. In fact, the 5G is the main telco application that will benefit from the virtualization technology and the automation & orchestration framework we are deploying.

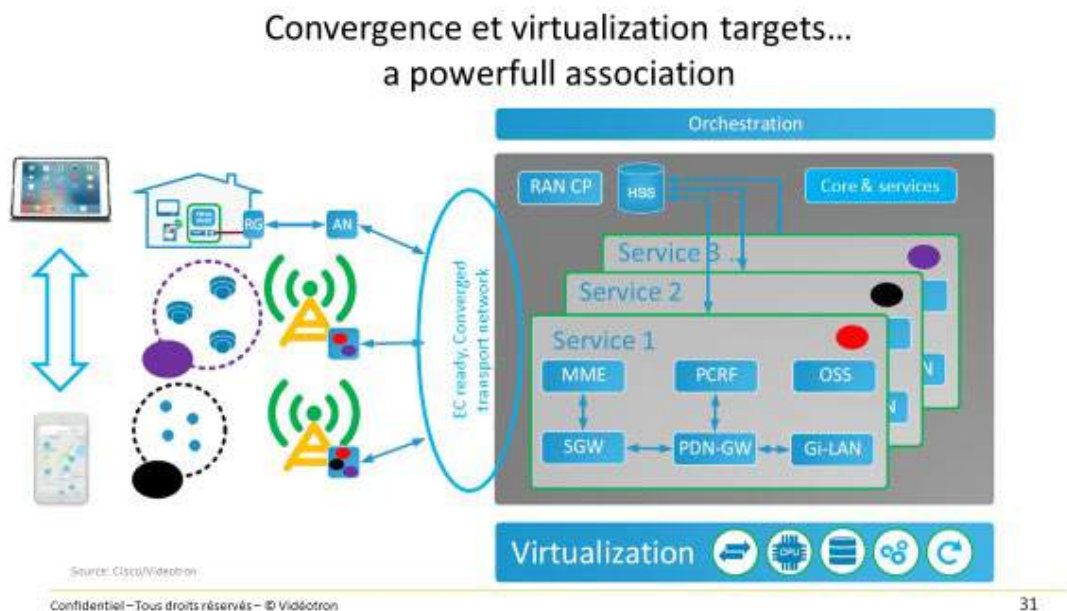
Through our automation and orchestration knowledge and process, we are targeting the creation of new values for our customers. The idea is to provide the right technological foundation for any new digital services. Any new digital services will require a specific part of the network and a specific location deployment (fog, edge, data center, public cloud). With network slicing, specific resources will have to be delivered for a specific B2B customer, and specific applications will be deployed near the customer to improve real-time experience. These use cases require a virtualized infrastructure and an automation & orchestration framework as shown in Figure 25.

One of the most challenging aspects is related to the maturity of the different virtual network function (VNF) partners to adequately run their VNF in an open virtualized ecosystem. They provide an entire virtualized stack that meet any operators needs in term of performance, scalability and service availability, but they require more investments than expected. Based on Gartner analysis and the open standard of the



industry (ETSI, TM Forum), we concluded that we need to limit the implication of such vendors in the value chain and spend more investments and time in the automation & orchestration area. The idea is to integrate properly any NFV stack at the OSS (our ecosystem for automation, orchestration, resource inventory, security tools and monitoring) level. That way, we let the NFV vendors be good at what they are good at, and we integrate their NFV stack to our OSS & BSS catalogue.

All the efforts done so far has created a new digital culture across the engineering team. The industry and our IT team has provided the right framework, knowledge and process to proceed to our NFV deployment. The future is to make sure our foundation will evolve the right way to host the 5G core, to enable new digital services area such as 5G network slicing and cloud gaming.



**Figure 25 – Convergence and Virtualization**




## 4.5. Telecom Argentina

### 4.5.1. Introduction

Telecom is the result of merging two former companies, one a mobile operator and one a cable operator. It has four main access network technologies that were evolving in different ways and then came together. The question is, how should Telecom evolve these networks in order to meet the capabilities and services that they will have in the future, while trying to choose the best of these different technologies, reusing the infrastructure, evolving the future access services, increasing the customer experience and avoiding duplicating investment due to the overlapping of different access network technologies in a same area.


Figure 26 shows Telecom's four access technologies, overlapping each other, in the same area: HFC (DOCSIS), fiber to the cabinet (FTTC) supporting xDSL, fiber to the home (FTTH) and 4G.



**Figure 26 – Telecom Access Network**

Figure 27 shows the deployment topology for each access in a cleaner fashion. Each of these networks has different capacities, characteristics and connection flexibility. For example, the HFC network reaches an area of 1,000 households with a couple of fiber optics (FO) from the hub to the fiber node, while the FTTC network reaches an area of 400 households with one FO. Even when the first one can provide more and better services, the second one provides very valuable growth potential for example with its underground wiring.

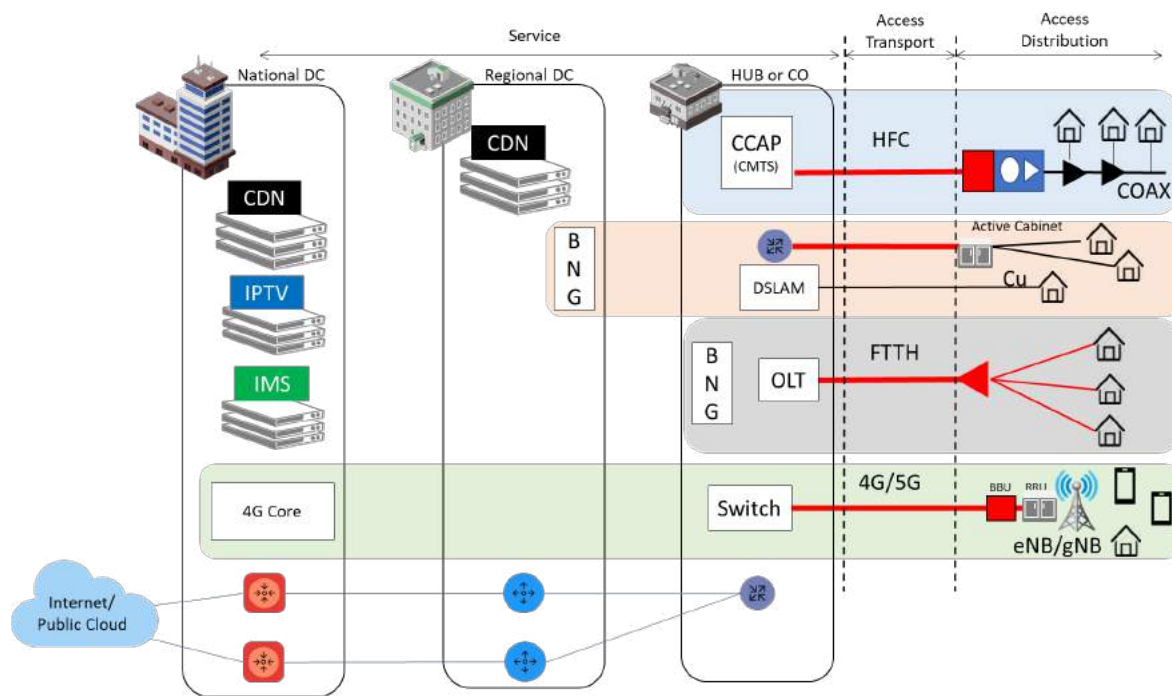
On the other hand, the 4G/5G wireless access network actually uses technology that could provide not only mobile services but also fixed services (especially on 5G). It also has a mobile backhaul over dark fiber to reach the base station from the central office. Finally, during the last few years, a new FTTH-based network has been implemented in certain areas and will be our target access network in long term.



**Figure 27 – Telecom Access Technologies**

The goal must be to find a **target access network** to provide the capacity and attributes for the future services, as well as a **converged access transport** to distribute and connect each access with the different core services in the data centers (those being national, regional, or even at the edges of the network).

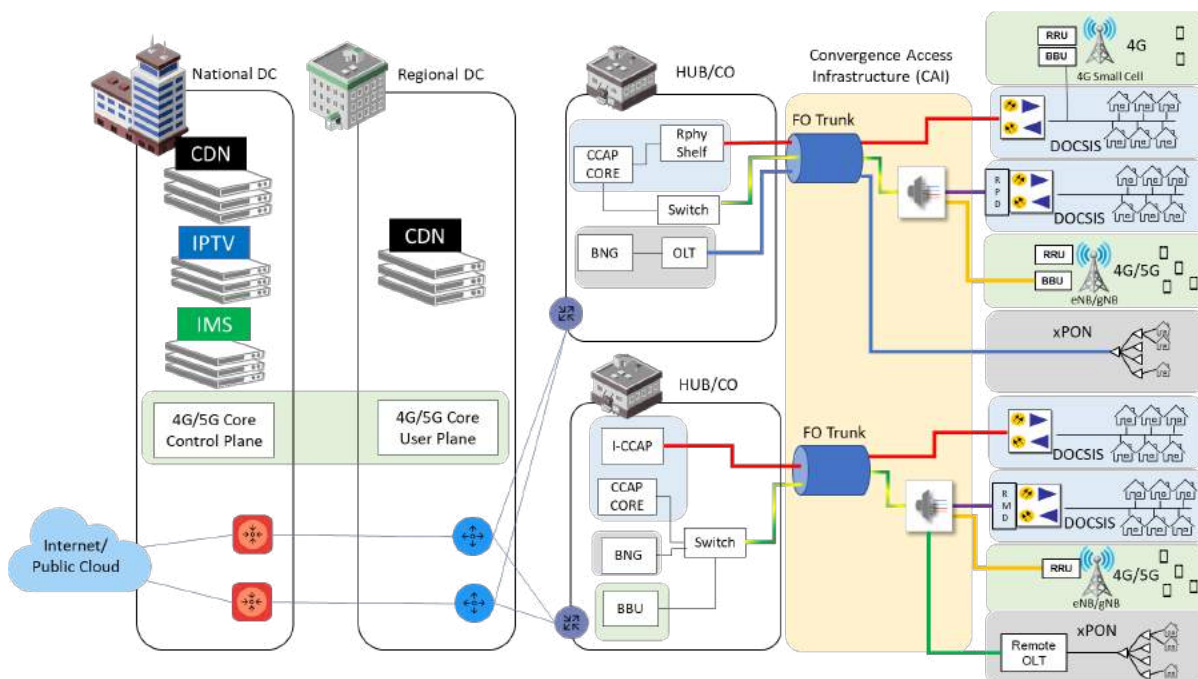
Figure 28 shows the complete picture of the four access technologies combining the access distribution, access transport network and the services that could be provided locally from the hub or central office (CO) or it could extend through different data centers (DC), like 4G in the national DC.



**Figure 28 – Telecom Access Network Prior to Convergence**

#### **4.5.2. Convergence Access Network and Convergence Access Infrastructure**

After several studies that Telecom did during 2019, based on different requests for information (RFIs), consultancy processes, and internal analysis, we found different approaches for Distribution Access Networks and Transport Access Networks. Figure 29 shows the new picture with convergence access network and transport infrastructure.



**Figure 29 – Access Convergence**

#### **4.5.3. Fixed Access Network**

For fixed access services, there is not just one recipe. Depending on the homes passed density of the area, take rates, product roadmap, and existing technologies deployed in those areas, there are different criteria on how to evolve the access network.

For instance, in areas where there are more than 50 homes passed per block, and take rates that are greater than 25%, and if the area has DOCSIS and DSL technologies, the analysis shows that the fixed access technologies evolution will first evolve on HFC capacity up to DOCSIS 3.1 on 1 GHz or 1.2 GHz spectrum and will then evolve towards FTTH overlay (GPON and XGSPON).

That means that Telecom will leverage the existing HFC/FTTC infrastructure to build an overlay network and it will smoothly move DOCSIS subscribers to xPON technologies. It doesn't mean that we are going to shut down the HFC network. Even in 10 years we are going to have our HFC network.

FTTH overlay doesn't mean a FULL FTTH migration. What the strategy of FTTH overlay means is that the target fixed access network should be FTTH at the end of the road. In our case, the period of analysis was 10 years, and just a percentage of the customers will be migrated to GPON and this will be a function of different factors such as product tier, churn, competition, and DOCSIS offload necessities. That strategy shows a smaller TCO compared to keeping HFC technology while evolving it towards DOCSIS 4.0. Two key points for that are:

- Leverage the existing fiber infrastructure from both former companies using a converged access transport network. Most of the HFC plant has aerial fiber and it could be extended if needed for optical distribution network (ODN), and there are underground fibers and/or fiber ducts from the former FTTC network. This strategy should lead us to an underground fiber trunk – whenever possible – towards field nodes, even active or passive.

- Home network services must be completely an IP environment. CMs must be replaced by optical network termination (ONTs) in case of DOCSIS to xPON customer migration. Home services must be the same for residential gateway (RGW) features, Wi-Fi capabilities, and monitoring. The most important thing is to protect the investment in set-top boxes (STBs). TV services are based on IPTV in the home network, which is something that Telecom already started to deploy (Telecom have analog clients, and the analog reclamation will be based on IPTV directly, so most of the home will be prepared to be migrated towards xPON during next years).

In areas with DOCSIS and digital subscriber line (DSL) technologies, but with lower households passed density and/or take rate, we found that the FTTH overlay was not so cost-effective in terms of total cost of ownership (TCO). However, DOCSIS evolution could support future demands, so we defined an option called HFC Target which moves xDSL customers towards DOCSIS. The risk of this option relies in ultra-high tier requirements driven by competition.

There are some older areas where the network is just asynchronous DSL (ADSL) technology or one-way HFC. These are in general suburban or rural regions with even lower households passed (HHP) density. Here, fixed wireless access (FWA) service was analyzed as the best choice to provide triple-play services.


The finding was that the best combination is to start the deployment with FWA, and then after some years, move to FTTH (assuming aerial fiber), and then reuse the FWA deployed capacity for mobile services. This alternative de-risked the investment of building a fixed network in areas where the take rate could not be properly estimated.

#### **4.5.4. Small Cell Backhaul over DOCSIS**

As depicted in Figure 29, there are small cell deployments over DOCSIS. The main use case here is to deploy pico-cells over HFC to improve the coverage service and provide better service quality to small, mid-size and large enterprises mainly inside offices, rooms, basements or to cover particular public indoor spaces (like shopping malls). The service must provide 4G and 3G connectivity.

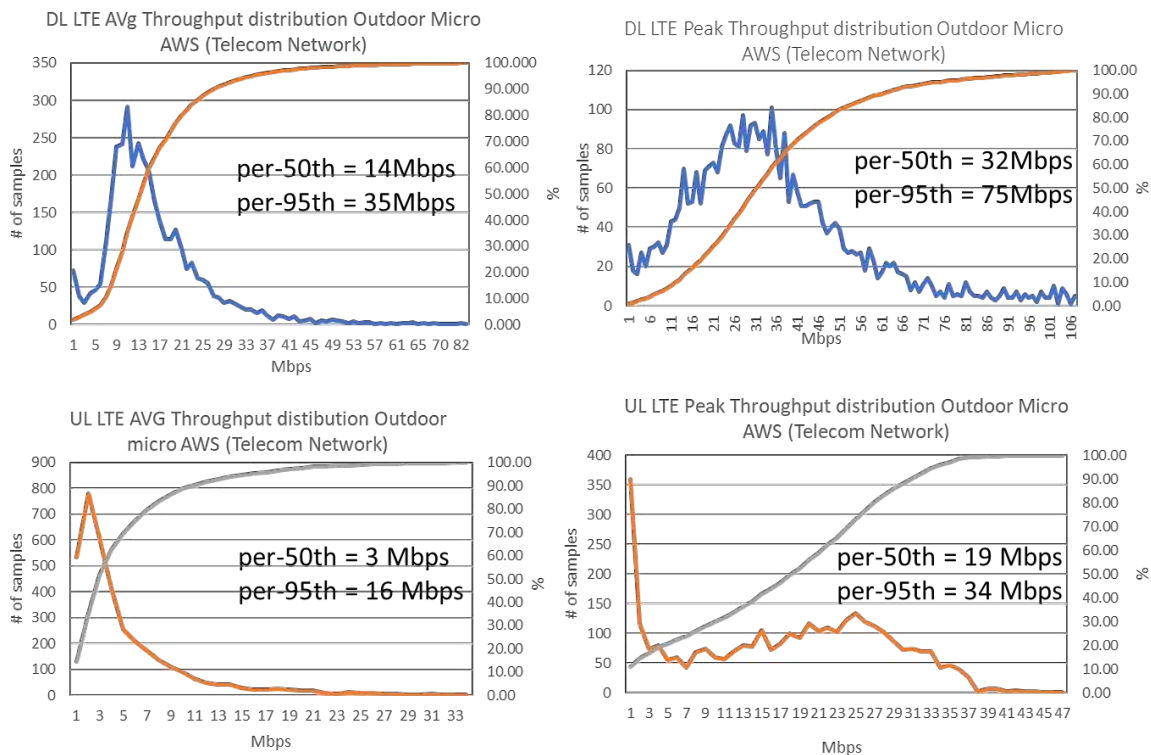
Figure 30 shows a simplified architecture that provides 4G services over DOCSIS backhaul. The HeNB (Home eNodeB is the name used in 3GPP to refer to small cells in 4G) is connected to the 4G EPC using DOCSIS as backhaul. The BBU is connected through a CM and establishes a secure connection with IPsec tunnel, that goes through the CMTS, the IP backbone (IPBB), and it is then terminated at the SeGW (Security Gateway).

Optionally, the architecture considers adding a HeBGW (Home eNodeB Gateway) that works as a concentrator of S1-U (i.e., user plane) and S1-MME (i.e., control plane) interface when there is a big amount of HeNB in order to avoid overload, especially in the MME. Finally, the HeMS (Home eNodeB Management) is the component that is in charge of the management of HeNB.



**Figure 30 – 4G Small Cells Backhaul over DOCSIS Architecture**

The following graphs show the distribution of average and peak throughput of LTE microcells for 2T2R<sup>3</sup> microcells during the peak hour, for downlink (DL) and uplink (UL).



**Figure 31 – LTE outdoor microcells Throughput**

For DL direction, LTE backhaul over DOCSIS is not a big concern in terms of peak traffic. Nowadays we have products of 100 Mbps in downstream (DS) for residential subscribers, then a small cell could be considered as another CM in the DOCSIS DS service group (SG). However, a small cell requires an average DS throughput of 35 Mbps, while a regular residential subscriber requires an average throughput of at least 10 times lower than that of a small cell. Therefore, sharing the DS SG between small cell deployment and residential subscribers must be done in a planned way.

<sup>3</sup> Measurements were taken in June 2018 – 2T2R Bandwidth = 15 MHz at AWS band.

From the upstream (US) point of view, assuming 70% utilization for capacity planning, 42 MHz split HFC networks provide approximately 47 Mbps (three 6.4 MHz channels with 64 QAM @70%) or 76 Mbps (four 6.4 MHz channels @ 64QAM @70%). On the other hand, the graph above shows that meeting the 95% of the times with the needs of UL LTE peak traffic demands a DOCSIS US SG capacity of at least 34 Mbps. Then in this configuration, it could be a challenge to use DOCSIS as backhaul sharing the resources with other residential CMs in the same service group. Hence, is highly recommended move the US split to the mid-split of 85 MHz and deploy D3.1 in US for more capacity, if the small cell is to share the same SG with residential subscribers.

Regarding 5G, in terms of capacity, if DOCSIS is considered as a backhaul solution, then mid-split and DOCSIS 3.1 in the DOCSIS US is a must, whilst in DS SG, the number of SC-QAM/OFDM blocks should be increased. Tests conducted in Telecom Lab with 5G gNodeB (gNB) showed DL peak rates of 1.5 Gbps and 130 Mbps in UL.

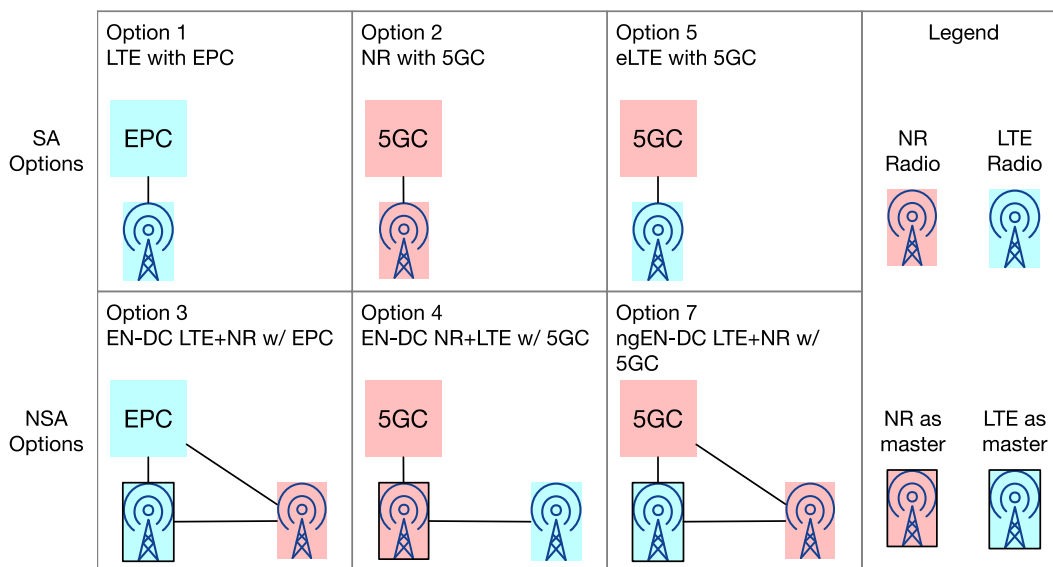
Another point to consider in backhaul over DOCSIS is the latency added to the end-to-end service. Round trip time (RTT) latency values that were obtained during the trials were in the range of 30 to 50 msec. In order to improve those values, the Low Latency Xhaul (LLX) over DOCSIS technology [6] is required (see Section 5.2).

#### **4.5.5. Evolution Towards 5G**

Deploying previous generations of mobile technologies, i.e., 2G, 3G, and 4G, involved abrupt changes, where, in addition to radio access, a complete change of the core was required to support the next generation. In stark contrast, 5G supports deployment alternatives that can leverage part of the 4G infrastructure, thus easing the inception of this technology.

The 3GPP defines a set of standard deployment options as shown in Figure 32. Regarding the evolution of mobile services from 4G to 5G, the first step is to deploy a 5G non-standalone (NSA) option 3x architecture [10]. Option 3x is one of the NSA options where an improved 4G EPC (EPC+) can be used to connect the RAN, which is composed of eNBs as master nodes as well as gNBs. Option 3x architecture gives Telecom a fast time-to-market and the possibility to provide enhanced mobile broadband (eMBB) services (for instance for FWA) in certain areas, as a complement to other fixed-wired services as well as 5G mobile services. For that, new gNBs are collocated next to the 4G's eNBs in areas where the service is required.



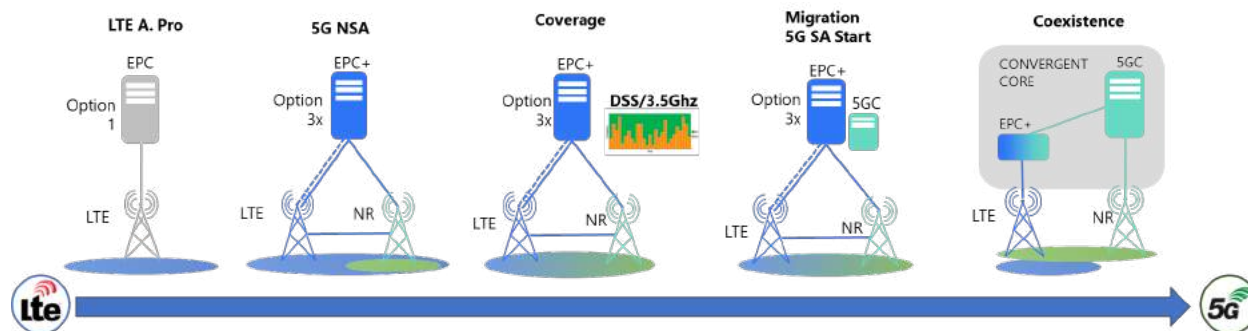


**Figure 32 – 3GPP Deployment Options**

Some upgrades are needed in the core to evolve the 4G EPC (EPC+) in order to support the NSA capabilities. At the same time, to improve the capacity and improve the traffic distribution, control/user plane separation (CUPS) architecture is implemented that splits the control plane in the national DC and user plane in the regional DC.

This will be followed by a coverage strategy that will be given depending on the commercial demand that we have, but also on the spectrum available. Here, there are several options. For example, using the spectrum in 3.5 GHz which requires regulation definitions in Argentina. Other options may be to make use of the existing bands (“refarming”). The problem is that it would take capacity out from 4G services.

However, a technique called Dynamic Spectrum Sharing (DSS) began to be developed. DSS allows dynamic use of the spectrum by arranging resources for 4G and 5G terminals at different portions of time and frequency subcarriers. It’s like doing an on-demand refarming. This is not really going to generate much more capacity than what we have today for 4G, because, at the end of the day, the amount of bandwidth is still the limit, but it will give the possibility that 5G terminals can camp into a 5G service area. Figure 33 shows a potential roadmap, from the initial state in LTE Advanced Pro, then 5G NSA core plus new generation radio deployment and coverage expansion.



**Figure 33 – Possible Evolution Toward 5G**



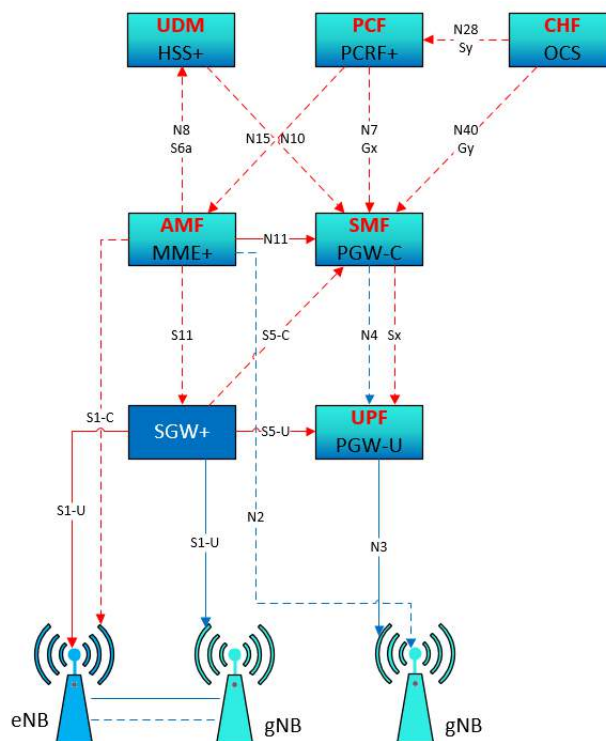
To offer the other 5G attributes, aka ultra-reliable and low-latency communications (URLLC) and massive machine type communications services, a new generation of radios (NR) are necessary but not enough. A full fledged 5G core using a service-based architecture is mandatory. This means beginning a process of core evolution and integration to include the new 5G standalone (SA) components in place.

Finally, we will have a converged 5G core with embedded capabilities that support 4G services since 4G will remain for many years still in our mobile access network. The first step is moving from LTE to option 3X NSA. The next step is to introduce 5G SA. In the industry, there are not much interest in option 7 or option 4. In theory, an operator could simply introduce a new 5G SA option 2 to connect new 5G UEs, then to operate 2 different cores, the “legacy” 4G and 5G NSA and the new one 5G SA.

However, while we introduce a 5G SA, we must guarantee the service continuity between LTE, 5G NSA option 3x and 5G SA option 2. Use 5G SA option 2 to connect 4G UE is another challenge because this migration strategy would require upgrading the eNB to gNB, while also would likely need to acquire new licenses in the new 5G SA option 2 to include the legacies 4G UE. This migration strategy requires more cost without any extra benefit.

What Figure 34 depicts is a combination where we keep the components of 5G NSA supporting 4G UE and 5G UE that works in option 3x, combined with 5G-SA components that connects the UE in option 2.

What we have is an evolution of our today 4G core including LTE+5GNSA 3x capabilities and finally we add 5G SA core components but all in the same “convergent core” solution. That allows us to keep the investment that we did for 4G and 5G NSA, to introduce the 5G SA, providing service continuity and giving the benefit to operate just one “platform.”



**Figure 34 – Convergent 5G Core**

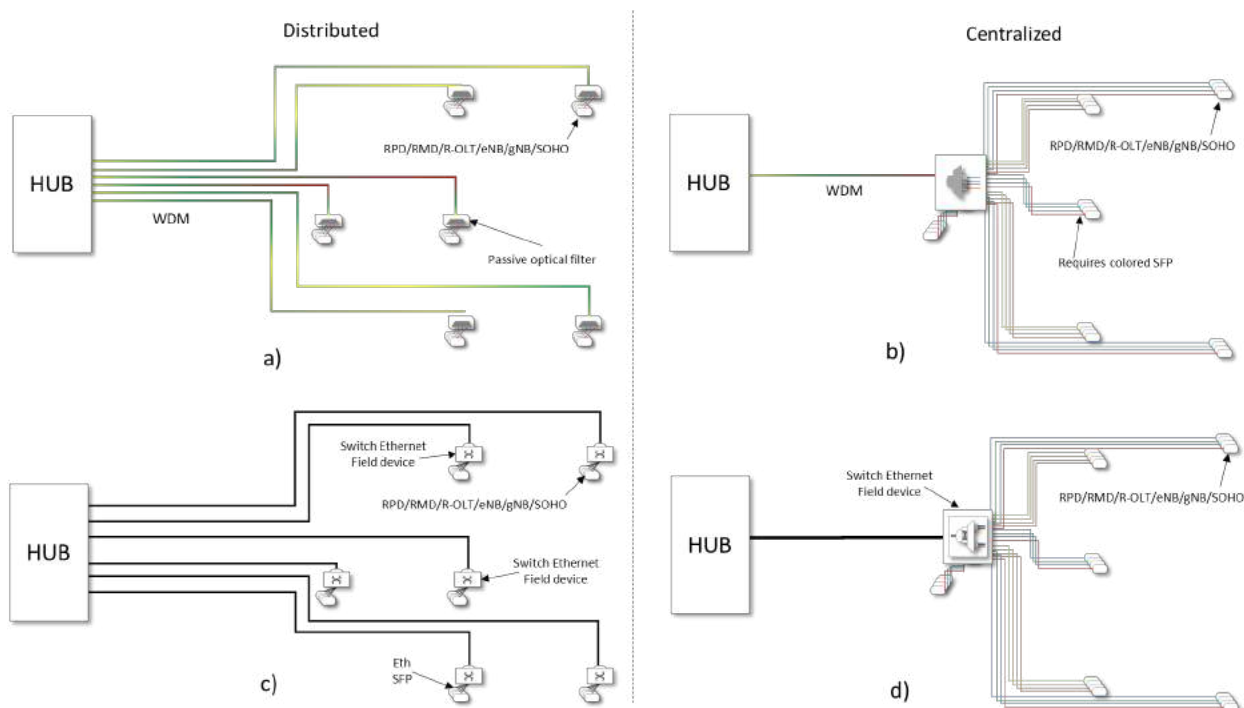
#### 4.5.6. Convergence Access Infrastructure

Figure 29 showed access distribution areas that are connected towards the hub and CO with a Convergence Access Infrastructure (CAI). That means that the same fiber trunks are used to support different services such as RPD, RMD, R-OLT, eNB, SMB CPEs, and to connect the access distribution area.

One option is to carry those services directly on dark fiber from the hub or CO to RPD, RMD, R-OLT, eNB, etc. In this scenario, each service uses one or two dark FO.

Another option could be multiplexing several services in one fiber and field devices to demultiplex or disaggregate the services. There are different techniques in this case: with passive architecture using wavelength division multiplexing (WDM), or with active architecture with Ethernet aggregation switches, and/or xPON technologies.

In Figure 35, (a) and (b) show passive WDM transport architecture. They could be distributed or centralized architectures respectively. In these cases, the field devices are passive optical filter and the services (RPD, R-OLT, eNB, etc.) requires a colored small form-factor pluggable (SFP). Figure 35 (c) and (d) depict an active Ethernet aggregation switch in the field, where (c) shows distributed architecture, and (d) shows centralized architecture.



**Figure 35 – Access Transport Technologies**

The realistic criteria to build the access transport network is based which part of the former companies' networks that is more appropriate. The analysis is based on using the part of the network that is convenient for Telecom due to its strategic location and getting the fiber vacancy to support the technology and the future service requirements. The goal is to protect the existing investment, while reducing the TCO and time-to-deployment, reducing the time-to-market, and avoiding delay that external plant construction could generate.

In the scenarios that Telecom has analyzed, in general, the WDM solutions are more cost effective than Ethernet aggregation switches solutions. The Ethernet aggregation switches solution requires a bigger initial investment. This is the reason why the rest of the analysis was done comparing WDM with dark fiber from a switch in the hub directly to the services such as RPD, RMD, etc., without any multiplexing mechanism in the fiber (dark fiber). The use of WDM techniques is the result of different considerations:

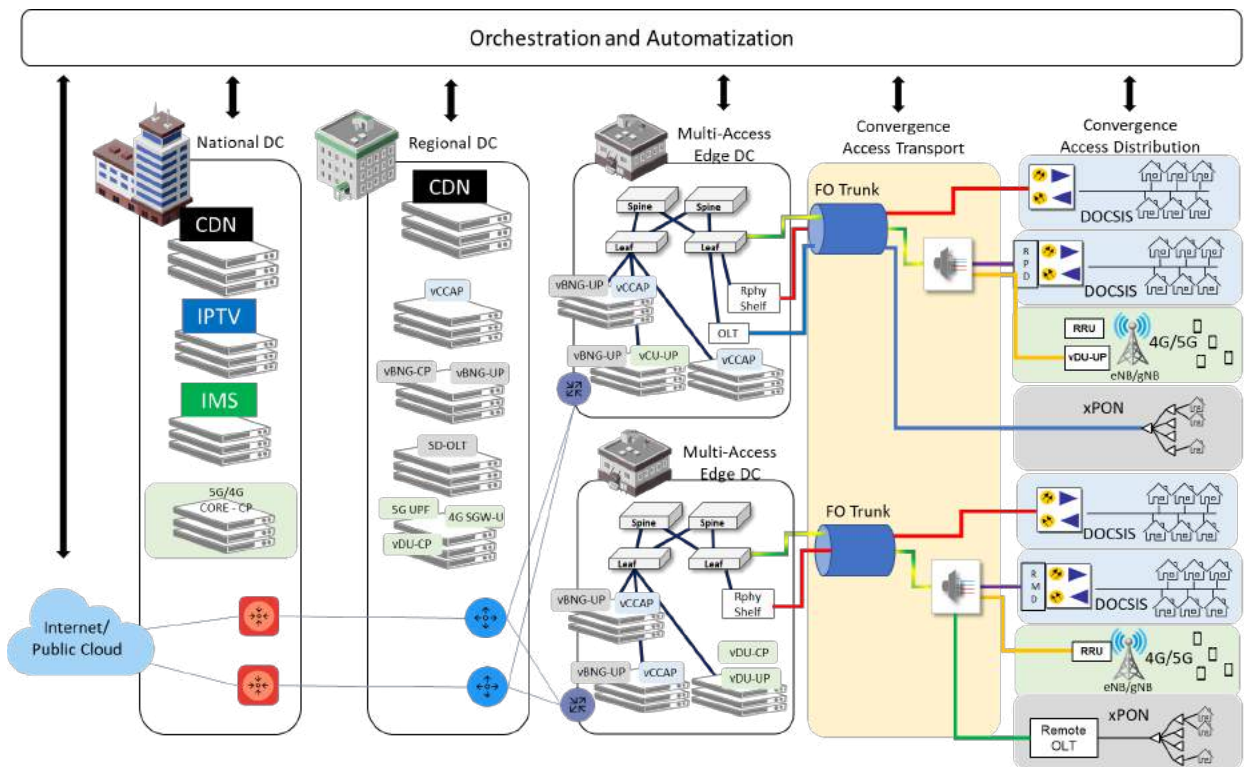
- Not every technology is suitable to be carried in WDM systems, as is the case of xPON over all passive external plant, referred to as the optical distribution network (ODN).
- The tradeoff between reusing the existing FO with WDM techniques and install new fiber trunks will depend on the distance of the trunk and the amount of the services to be carried (DOCSIS, mobile backhaul/midhaul, business services, remote OLTs). WDM could accelerate the time-to-market because it requires less deploying time and less labor force than deploy new fiber trunks.
- As an alternative to new fiber cable installation, neither of them represents significant TCO compared with fixed access infrastructure (CPEs, ODN, etc.), but WDM accelerates time-to-market and does not require labor force for cable landing.
- WDM introduction can be cost effective when the plant evolution is distributed in time, i.e. an area with many HFC nodes that has to be evolved with FTTH overlay at one node per year.
- In the FTTH overlay and HFC scenarios, as the external plant requires active equipment in outside plant, the use of remote optical line termination (rOLT) for residential and some small office home office (SOHO) customers, and direct fiber for bigger customers, the use of WDM optimizes the reuse of existing fiber infrastructure without regard for the former company's origin.

#### **4.5.7. Virtualization and Cloudification**

Access services have dedicated infrastructure for CMTSs, OLTs, and broadband network gateways (BNGs). With the network virtualization, or network cloudification, there are some components of those equipment that will be virtualized. Several functions that are in a CMTS are also in the BNG. For instance, subscriber management, DHCP, IP/MPLS forwarding, routing protocol, and others. Of course, there are certain functions that are only characteristics of technology with CMTS, BNGs and OLTs, but those also could be virtualized and there are other functions that will keep as physical network functions.

Nowadays most of the developments in the industry are moving from a legacy concept of virtualization to the new one, that is “cloudification.” The “legacy” way to virtualize a network function has been changed with the new cloud native paradigm. This means rebuilding the system based in open-source software components and in microservice oriented architecture, where those microservices are containerized and orchestrated dynamically.

Microservice architecture means that the system is divided into small applications that could be developed and scale independently of each other, improving and accelerating the agility, maintenance and development of new capabilities. Containerization provides a light way to virtualize each microservice or process, it could be packaged in isolated way, which provides an easy way to be reproduced and deployed. Finally, the orchestration and automation manage and schedule the resource utilization, providing mechanism to simplify the deployments and scaling of containerized applications.



**Figure 36 – Virtual or Cloud Access Services**

Those virtual network functions or cloud native functions will share almost the similar physical and logical infrastructure. That is a big advantage if we think in FTTH overlay architecture where we could reuse the same infrastructure to deploy capacity for CMTS, BNG and OLTs. Furthermore, if we think in a DOCSIS to FTTH migration, we could discuss with the vendors to move licenses cost from a virtualized CMTS to a virtualized BNG, for example, to protect investments. Of course, we would not replace the CMTS and BNG that we have today with virtualized versions. However, when capacity expansion is needed, we will deploy the new virtualized CMTSs or BNGs instead of the traditional physical versions.

Virtualization would also enable us to move some of the functions to a more centralized location, such as a regional data center (DC), to have a better scale, and leave in the hub/CO, or the new edge DC, or far edge DC, just the functions required to guarantee latency or to keep capacity closer to the end user. Regarding the optimization of the location of different components, the CUPS architecture applies too. The BNG is a typical use case of CUPS where the control plane (CP) can be in a regional or national DC, and the user plane (UP) is distributed across edge DCs. This substantially improves the management of BNG infrastructure.

## **4.6. Vodafone**

### **4.6.1. Background and Context**

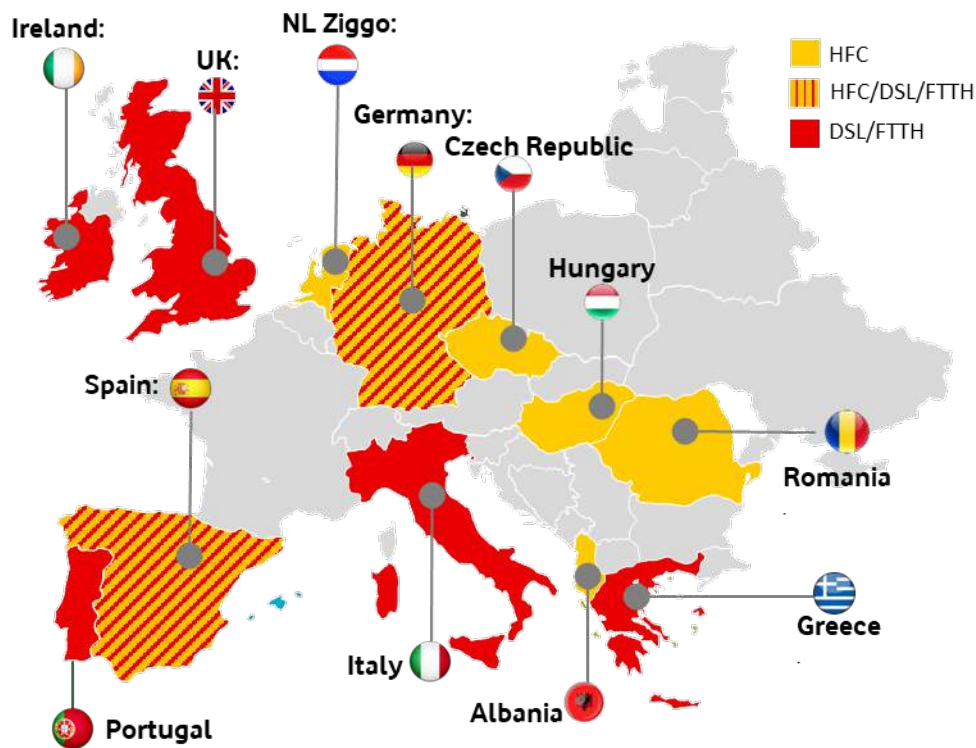
Vodafone started as a mobile operator in the UK in 1985 and subsequently expanded its mobile operations into other countries around the globe via a combination of owned networks and partner markets to reach over 300 million customers. Vodafone has mobile operations in 22 countries and partners with mobile networks in 42 more.

Vodafone has grown its fixed line business via a combination of acquisition, self-build and joint ventures. Examples of fixed network acquisitions include Arco in Germany, Tele2 in Italy and Spain, C&W in the UK and Hellas Online in Greece. Vodafone then became the largest local loop unbundler in Europe using its own ADSL and single-pair high-speed DSL (SHDSL) equipment on rented copper lines from incumbents. It also used sub-loop unbundling in Italy and Germany to deploy its own very high-speed DSL (VDSL) equipment in street cabinets.

Vodafone has been deploying FTTH for over a decade and, for example, has now fiber-passed over half the population in Portugal. It has also expanded FTTH via partnerships and joint ventures (JV) such as in Ireland where SIRO is a FTTH JV between Vodafone and ESB, the electricity utility.

In terms of cable, Vodafone acquired KDG in Germany, ONO in Spain and more recently the Liberty Global cable networks in Germany, Czech, Hungary and Romania. Vodafone also acquired ABcom in Albania. This has resulted in Vodafone becoming the largest broadband operator in Europe and one of the largest cable operators globally.

Vodafone provides fixed broadband in 17 countries. As of 30 June 2020, Vodafone Group had 27 million fixed broadband customers and 22 million TV customers, including the customers in Vodafone's joint ventures and associates. Vodafone's European fixed broadband technology presence is illustrated in Figure 37.



**Figure 37 – Vodafone's European Presence**

Outside of the European footprint above, Vodafone also has operations including fixed broadband in Turkey and a number of AMAP (Africa, Middle East, Asia-Pacific) countries too including South Africa and Ghana.

Vodafone is now a converged operator whose roots are in mobile. Hence in comparison to many US and Canadian cable operators, Vodafone has significant mobile spectrum and network assets. It therefore has a slightly different approach to convergence in that it is not an MVNO and it does not seek to use the CBRS spectrum. Also, whilst Vodafone has experience of deploying Wi-Fi hot spots in some markets, it is not highly dependent on having a vast number to offer “wireless untethered access” outside of the home.

Vodafone has 2G, 3G, 4G and 5G networks. Mobile network evolution will significantly grow 5G coverage, capacity and functionality over the next few years in addition to other mobile network capabilities such as narrowband IoT (NB-IoT). Mobile innovation will continue with developments such as Open-RAN and Crowd-Cell. There are also opportunities to couple such mobile capabilities with our fixed network assets for convergence at the network and service layers. The following sections give an overview of Vodafone's approach to convergence.

#### **4.6.2. Motivation for Convergence**

It has been over a decade since some service providers offered “triple-play” service bundles of fixed voice, broadband and video. Convergence refers to combining both fixed and mobile capabilities resulted in “quad-play” with the addition of mobile voice/data to the bundle. This has become increasingly common across Europe where some markets such as Spain are highly converged with multiple operators offering converged service bundles to customers.

There are three main commercial benefits that have motivated Vodafone's move to become a converged operator:

### 1. *Reduced Churn*

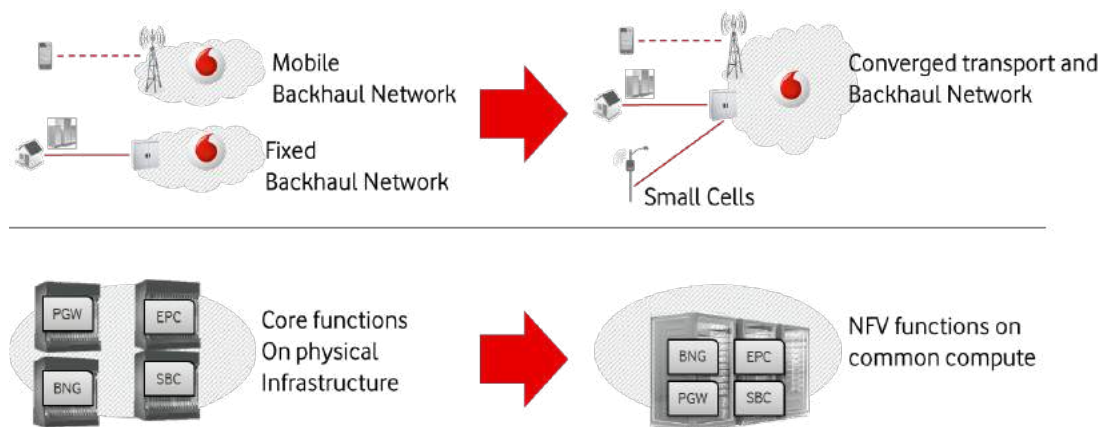
Customers buying a bundle of converged fixed plus mobile services tend to be “stickier” and less likely to churn. They are often incentivized by getting a discount on the bundle (compared to procuring each service individually), and/or more included benefits (e.g. extra data, additional mobile lines, higher speed).

### 2. *ARPU Improvement*

When a service provider becomes a converged services provider for the first time then they have the opportunity to cross sell fixed services to mobile customers and vice versa.

### 3. *Cost Reduction*

Convergence at the network level enables savings compared to operating disparate fixed and mobile networks. Hence, this is usually one of the first activities to drive synergy savings following M&A activity between fixed and mobile operators. The fixed and mobile services can then share common backhaul and core transport networks as well as data centers and server capacity, which is increasingly important as we move to software-defined virtualized networks. This is illustrated in Figure 38.



**Figure 38 – Infrastructure Convergence**

The simplest form of convergence is on the customer's bill – a single bill to span both fixed and mobile services. However, more sophisticated convergence approaches are feasible, convergence can take several forms from convergence at the service level through to convergence at the network infrastructure level. A simplified view of convergence in four different domains (or layers) is illustrated in Figure 39 for the Vodafone context.





**Figure 39 – Vodafone Four Layers of Convergence**

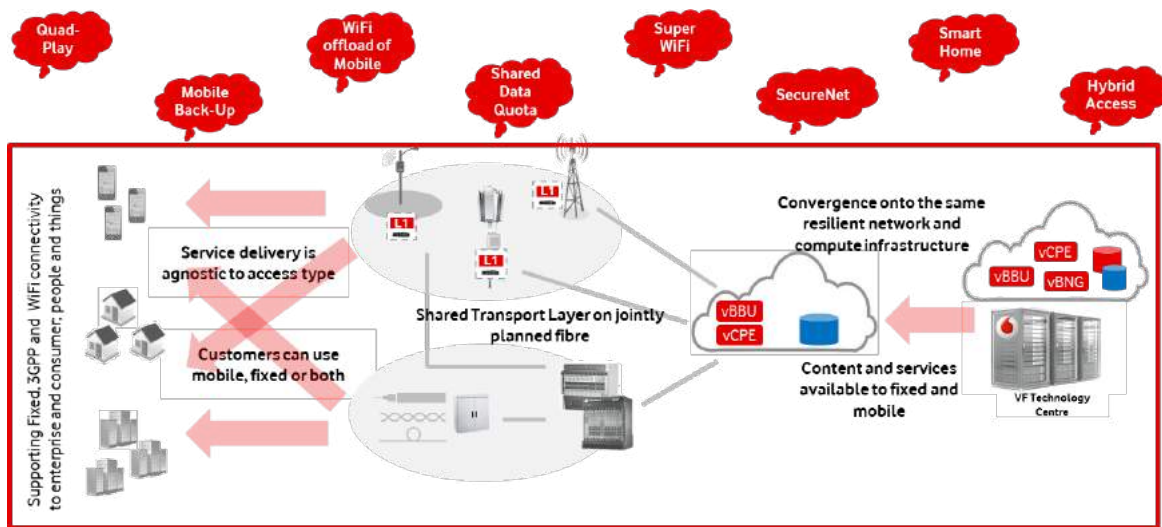
At the bottom layer we have network connectivity and infrastructure such as compute processing capacity in data centers. The next layer has the capabilities, systems and processes that turn the network and compute resources into a technology platform. This is where policy, analytics, operations processes and automation (SDN control and orchestration) reside. This layer interfaces to the Services and Application layer above it via standardized APIs (effectively a Network Exposure Layer) to create a “Network as a Platform” paradigm. Then any converged product and service development can focus on developing to the APIs (ideally TM Forum compliant) with worrying about details of the underlying network technologies and their associated idiosyncrasies.

Development of new products and services can then be more rapid and also more rapidly deployed across multiple markets (without repeating integration heavy lifting for each new geography). This benefits time to market for internal Vodafone product development teams but also facilitates engagement with third-party partners. Finally, the top layer is for Channels and Customer Service – effectively the touchpoints for the end user customers.

#### **4.6.3. Convergence Use Case Examples and Network Scenarios**

A converged network, as illustrated in Figure 40, allows a variety of converged services to be offered.





**Figure 40 – Convergence Use Cases**

The following sections discuss examples of converged services and network capabilities.

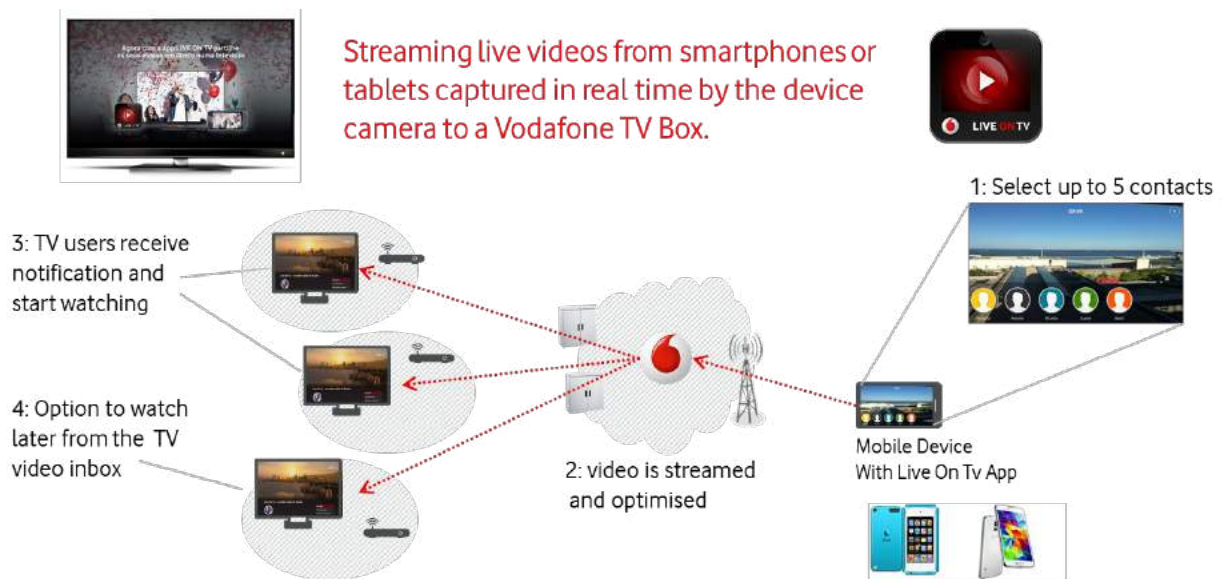
#### **4.6.3.1. Vodafone TV**

This enables a seamless, device agnostic TV experience in and out of the home which is personalized across all sources. It gives easy access to all content – integrating traditional/linear channels with on-demand, catch-up and the best OTT streaming services – with advanced search, including voice. The banding is shown in Figure 41.



**Figure 41 – Vodafone TV**

A further converged Vodafone TV variant that had previously been deployed in one of Vodafone's local markets is illustrated in Figure 42.




**Figure 42 – Local Streaming**

For example, if one of your children was playing a football game on a cold, wet Saturday morning, someone from the family could attend and use their smart phone camera to “video” the game which would be relayed to the TV where the rest of the family can watch it in warm, dry comfort. The process of making this happen is shown in Figure 42. There are now alternative ways to achieve similar capabilities since over-the-top (OTT) platforms (Facebook, YouTube, etc.) can offer live streaming capabilities and cast them to the TV too.

#### **4.6.3.2. Safety, Parental Control & Smart Home**

A range of “Smart Home” services is illustrated in Figure 43 are connected to the fixed broadband access line (via Wi-Fi, Zigbee etc.), but the control point is the customer’s mobile handset. Such services can provide customers with control in and out of the home, across all devices, fixed and mobile. Capabilities can include mobile control and notifications. Security capabilities can include alarms, locks, smoke and leak sensors, IP cameras. Smart energy can also be facilitated via remote mobile control of thermostats.




**Figure 43 – Smart Home**

#### **4.6.3.3. Hybrid Access (Fixed-Mobile Bonding)**

Hybrid access combines the throughput of a low-speed fixed broadband connection such as ADSL with mobile 4G/5G data using a hybrid router. This enables customers to burst to higher speeds beyond the capacity constraints of their fixed broadband connection. It can be used to provide an ‘always-on’ service i.e. a resilient service with mobile backup in case of fixed service failure.

Vodafone’s initial deployment of this capability, optimized for the small-medium enterprise market, integrated the bonding client into the consumer broadband CPE. This is shown in Figure 44. A variant is also feasible that uses a customer’s mobile handset to provide the mobile access connection to temporarily boost the customer’s fixed broadband access speed.



**Figure 44 – Hybrid Access**

Other variants of such technologies (including 3GPPs ATSSS) can also help to facilitate seamless roaming between mobile and Wi-Fi connectivity, preferably whilst maintaining the same IP address. This can ensure that the customer always has the best connectivity (based on both RF signal strength and network congestion). Such mobile/Wi-Fi roaming capabilities also have Enterprises use-cases e.g. in hotels, conference facilities etc.

#### **4.6.3.4. Always-On Service**

A precursor to the hybrid access approach illustrated above was to leverage the availability of both fixed and mobile access connectivity in a more manual way in order for the customer to be able to carry on accessing the Internet if their fixed broadband fails. This was deployed in a few countries including in Vodafone Greece.




**Figure 45 – Always-on Service**

#### **4.6.3.5. Fixed Wireless Access (FWA)**

FWA is a technology option as illustrated in Figure 46 is considered for low to medium speed broadband access for areas with no or poor fixed access. It can be used as a tactical deployment where wholesale fixed broadband costs are excessive. It is usually avoided for deployment in urban and sub-urban (Residential) areas with good next generation access (NGA) penetration. It is often a “last resort” technology because there can be high customer equipment costs (especially if outdoor hardware and professional installation is required).

There can also be significant capacity implications for the costly spectrum on the mobile network (e.g. somebody binge-watching a 4k video stream for a few hours. Hence such solutions are most commercially viable in areas underserved by high-speed fixed networks and with low customer penetration, or amongst specific customer segments (e.g. customers who move residences frequently etc.). The Vodafone “Gigacube” 4G/5G FWA product has been deployed in Germany, UK and Ireland.

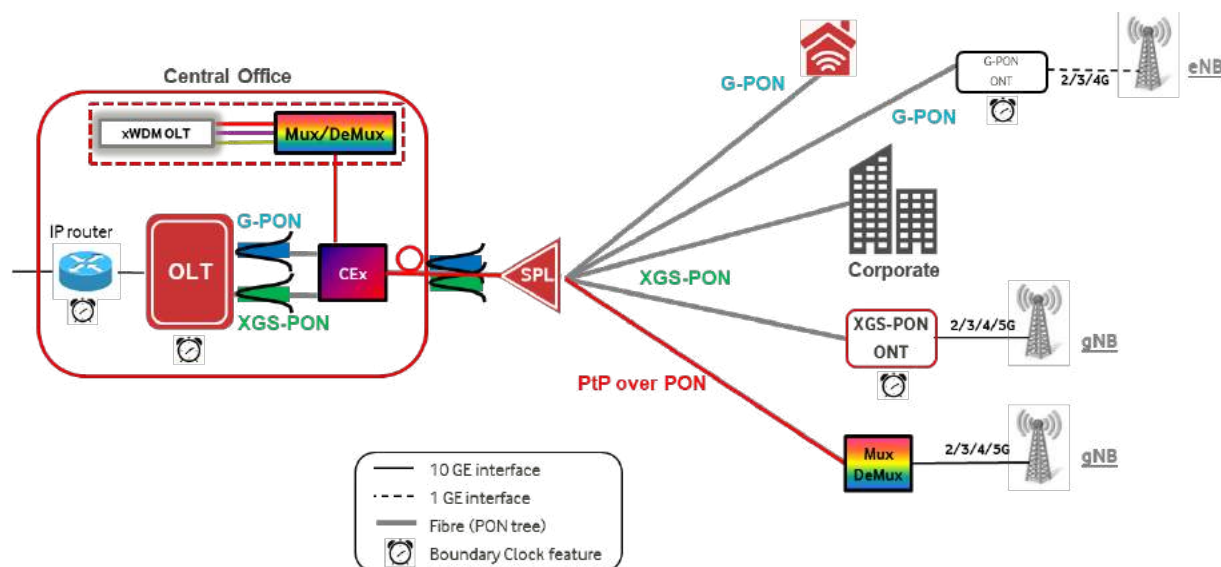


**Figure 46 – Fixed Wireless Access**

#### **4.6.3.6. Converged Access & Aggregation Network**

As NGA technologies have evolved, it becomes feasible to examine their potential to provide backhaul from mobile base-stations as an alternative to point-to-point fiber and leased lines. Vodafone had previously used GPON to connect a few base-stations in some countries. This approach then evolved to evaluate PON backhaul from 4G and 5G base-stations including the evolution from GPON to XGS-PON. XGS-PON products including the necessary Synch support (IEEE 1588v2 & SYNC-E) and have been

successfully proven. DWDM over PON is also a potential technology option which is under investigation. This is illustrated in Figure 47.



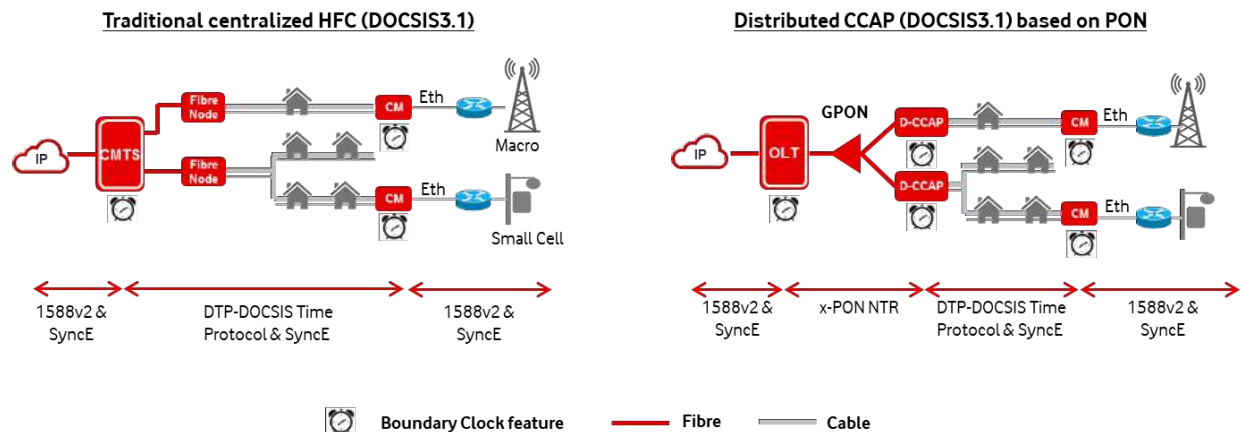
**Figure 47 – Converged Access with DWDM PON**

Thus far, mobile RAN developments such as SON (self-optimizing networks) and active antennas (also known as massive MIMO) have successfully increased the RAN capacity from macro-cell base-stations, which has deferred the need for large-scale deployment of small cells. Nevertheless, where required, Vodafone has deployed a number of external public small cells as well as indoor femto-cells for consumers and pico-cells for business customers to enhance indoor coverage.

Vodafone has previously trialed mobile backhaul over DOCSIS 3.0 in order to understand synchronization and bandwidth utilization requirements. The focus then shifted to examining mobile backhaul over DOCSIS 3.1. This is more challenging for macro-cells compared to using XGS-PON due to the lower bandwidth and potential implications for residential broadband users. However, there is a potential future role for using DOCSIS to provide backhaul from small cells, as and when their deployment may become necessary at any significant scale.

For mobile backhaul over DOCSIS, the phase and time requirements can be met due to the inclusion of the DOCSIS Time Protocol (DTP) in the DOCSIS 3.1 standard. QoS and mechanisms to manage the jitter are needed. Also, low latency (e.g. leveraging CableLabs Low Latency Xhaul – LLX) is needed to stay within a target latency of 5 ms (base station to backhaul aggregation). The various synchronization mechanisms are shown in Figure 48.

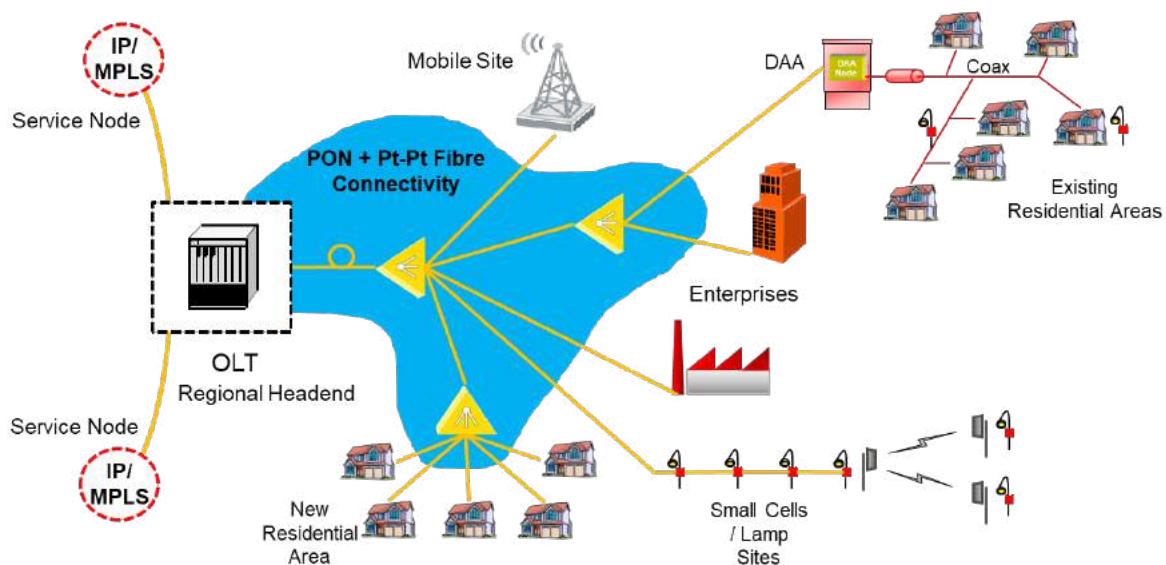




**Figure 48 – Network Timing with DOCSIS and PON**

Vodafone has been evolving its cable networks to a distributed architecture in some markets as we introduced DOCSIS 3.1. This requires fiber to be pushed deeper into the access network to connect to a remote PHY device (RPD) and remote MAC device (RMD) remote devices. This trend will continue as we evolve towards DOCSIS 4.0. Hence it creates the opportunity to have a single fiber access and aggregation network that can provide direct fiber connectivity (PON or point to point) to macro base-stations, enterprise customers, business parks and the cable network’s remote nodes.

As more bandwidth is required on such a “unified fiber access/aggregation” network, coherent optical technology has the potential to play a key role, especially for the backhaul from remote aggregation locations. This converged fiber access/aggregation network approach is illustrated in Figure 49.



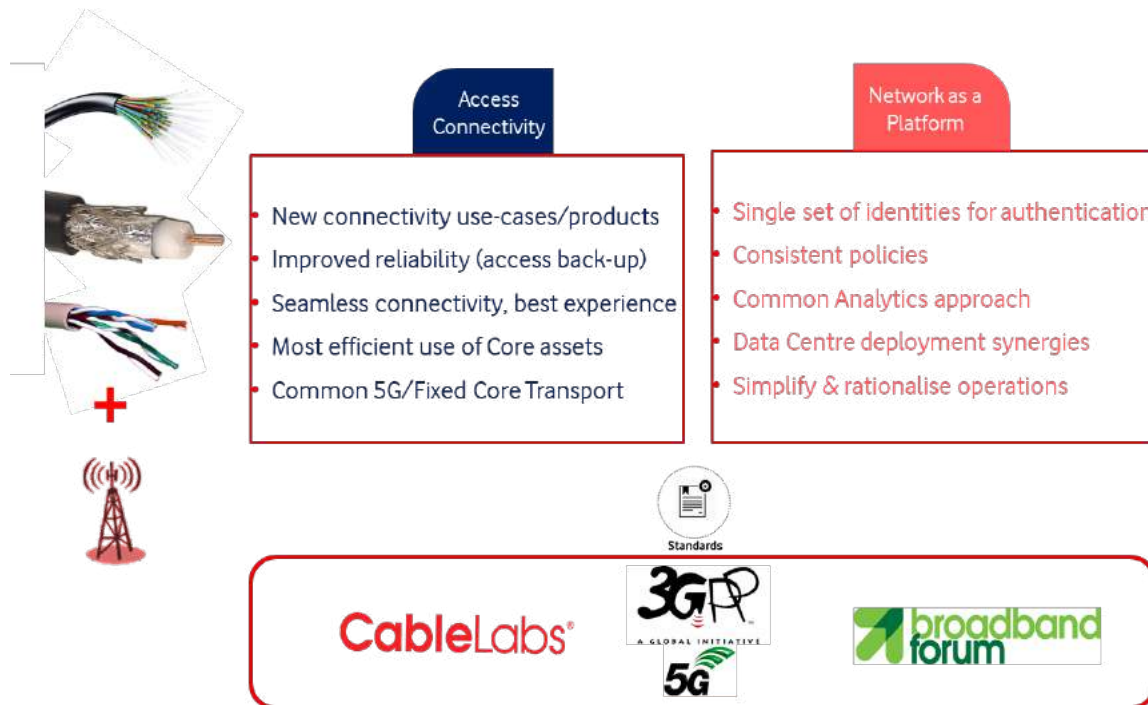
**Figure 49 – Unified Fiber Access/Aggregation Network**

A further enhancement to this access network is to introduce session steering or slicing from the access node in order to steer sessions to a user plane (software or hardware) that is appropriate in terms of location (latency) and scalability/cost to the particular session’s traffic. For example, low margin consumer traffic from OTT streaming could be steered to a scalable, cheap (and dumb) user plane whereas enterprise traffic

may be steered to a user plane where service chaining is to occur for value-added services like security (firewall, malware scrubbing etc.). In the convergence context, LTE mobile backhaul traffic could be steered to the Gi LAN.

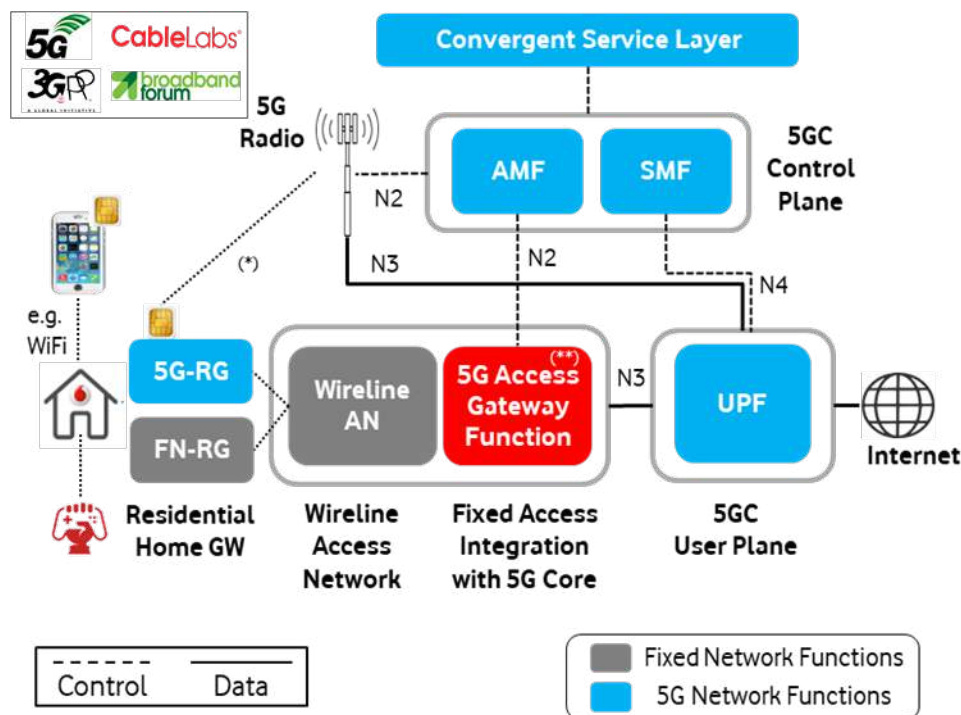
#### 4.6.3.7. **Broadband Access to Converged 5G Core**

There are a number of incentives for ensuring that the next generation 5G core network can support wireline access instead of treating it as “untrusted”, as was the case in previous 3GPP mobile core network standards. These are summarized in Figure 50.



**Figure 50 – Wireline Support in 5G**

The technical approach has been to develop an Access Gateway Function (AGF) that effectively translates from the wireline protocols to those used by 3GPP. 3GPP has been working with both CableLabs and the Broadband Forum on standards in this area. The network context of the AGF is shown in Figure 51.



**Figure 51 – Wireline Connection to 5GC with AGF**

Residential gateways (RGs) will also evolve to encompass functionality that enables enhanced capabilities. These new RGs are denoted “5G-RG” in the standards documents. The new “5G-RG” functionality in CPE will enable more dynamic session connectivity. A 5G device (UE) behind the 5G-RG can be treated as “trusted”. QoS requirements can then be signaled on a per application basis and multiple IP sessions can be added dynamically for a class of device. The 5G-RG is illustrated above. The “FN-RG” is the legacy fixed network RG that can be used to access the converged 5G core but will have less flexible/dynamic capabilities than the 5G-RG.

#### **4.6.3.8. Improving In-Home Mobile Voice Coverage**

Vodafone has extensive experience of using femto-cells plugged into the customer’s broadband router to improve in-home mobile voice coverage. A product known as “Sure Signal” which was a 3G femto-cell was used in significant numbers in Vodafone UK. During Covid-19 lockdown in 2020 it proved useful for some of Vodafone’s own staff where, for example, their home office was in the basement or modern building materials (especially glazing) meant their home was effectively a Faraday cage.

An alternative to femtocells for improving mobile voice coverage is to enable VoWiFi on the mobile network (assuming internal Wi-Fi coverage in the house or basement is better than the mobile coverage). This is supported and easy to set up on modern smart phones. This is shown in Figure 52.





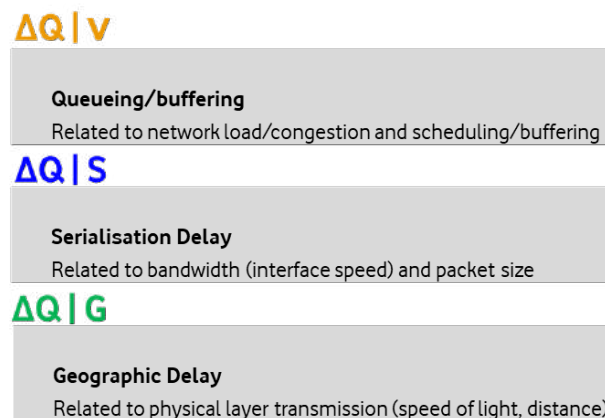
**Figure 52 – Voice over Wi-Fi**

#### **4.6.3.9. Converged Measurement & Analytics**

As both fixed and mobile access speeds have increased, latency (and other aspects of “broadband quality”) have become more important. 5G in particular was the first technology to use ultra-low latency in its marketing. Cable technology has also evolved to include Low Latency Xhaul (LLX) and Low Latency DOCSIS (LLD). In addition, CableLabs has worked on Low Latency Wi-Fi.

Vodafone’s recommended latency measurement method at the IP layer has been TWAMP (Two-Way Active Measurement Protocol). However, as we seek to further optimize customer experience and application performance we needed a more hi-fidelity technique for latency and performance measurement.

After extensive scouting and testing we selected the new Quality Attenuation technique, sometimes referred to as  $\Delta Q$  (pronounced Delta Q).  $\Delta Q$  emphasizes the gap between real performance and perfection (i.e. zero loss and delay) and has been standardized in BBF TR-452.1 with some early vendor capabilities available. The Quality Attenuation measurement and analysis approach can disaggregate a round trip time (RTT) into three constituent components as shown in Figure 53 in each direction (downstream and upstream), for a total of six components.



**Figure 53 – Quality Attenuation technique**

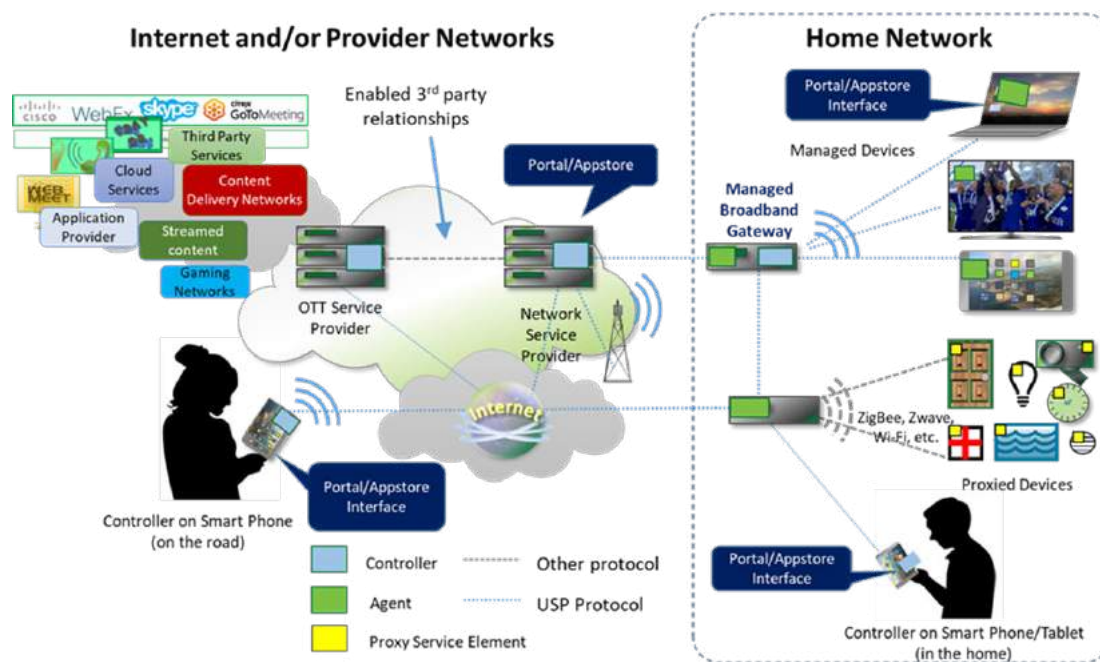
Vodafone has trialed the Quality Attenuation technique in four countries over a range of fixed and mobile access technologies including DOCSIS, 4G, GPON and VDSL. Different fixed and mobile access technologies behave differently at the physical layer and these manifests themselves as different packet latency and loss characteristics at the IP layer, which in turn determines the application outcome and customer experience.

50 Mbps on 4G is not the same as 50 Mbps on DOCSIS or FTTH. It is increasingly important to understand such issues as in the converged network we increasingly seek to use FWA, hybrid access to boost speed or to back-up fixed access with a failover to mobile access. Common tools for measurement and analytics across the mobile and fixed access elements of the converged network will become increasingly important as we seek to deliver our customer's applications in a seamless manner, irrespective of access technology.

#### 4.6.3.10. CPE Management

Vodafone has millions of Customer Premises Equipment (CPE) devices managed using the TR-069 protocol. Globally, there are now over a billion devices worldwide managed by hundreds of service providers using this protocol which is included as an option in CableLabs' eRouter standard. 3GPP and the Small Cell Forum have also developed data models to be used with TR-069 for the management of small cells.

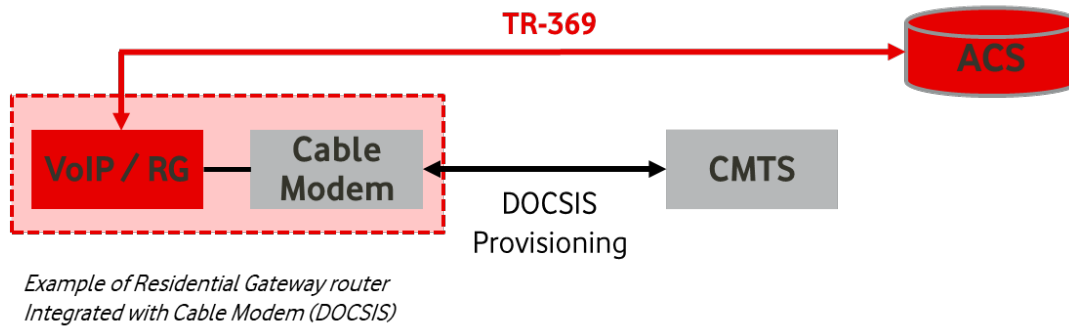
TR-369 is the successor to TR-069 and is a new standard known as the User Services Platform (USP). It has a number of improvements but the most notable is the ability to have multiple controllers. This facilitates new commercial models with external 3<sup>rd</sup> party partners for example the ability to outsourcing Wi-Fi optimization to an external partner with smart algorithms, AI and machine learning based in their own Cloud. They would only be able to access the Wi-Fi parameters in the Residential Gateway. An overview of user services platform (USP) as defined by TR-369 showing the multiple controller capability is illustrated in Figure 54.



**Figure 54 – USP TR-369 in Deployment**

In the convergence context, this could, for example, enable the fixed broadband experts to manage the residential gateway but the mobile experts to manage any connected small cells.

For an integrated CM and Residential Gateway (router), TR-369 can be used in conjunction with existing DOCSIS provisioning systems and processes as illustrated Figure 55.



**Figure 55 – DOCSIS GW with TR-369**

#### **4.6.4. Summary**

This section has presented how a range of technologies can be used to facilitate convergence at the network and service layers for a converged operator. It is important to note the key role that standards play in this evolution. In the past, standards have been siloed in either the fixed domain or the wireless domain. In an increasingly converged business, it is vital that CableLabs, 3GPP, Wi-Fi Alliance and Broadband Forum and other such organizations collaborate and cooperate in order to deliver effective solutions for network operators and service providers.

This is especially true for capabilities above the physical layer such as architecture frameworks, management protocols, data models, performance measurement techniques, telemetry and analytics. Alignment on such areas will be key in the new software-centric world of virtualized, distributed and disaggregated networks with a growing focus on automation leveraging SDN, AI and ML.

## 5. Convergence – Technologies

In the previous section, we covered business objectives and technology viewpoints of the cable-mobile operators. In this section, we recap some of the common technologies that have been highlighted to enable convergence between cable and mobile deployments.

### 5.1. DOCSIS Technology

Every wireless network is dependent on a wireline network. Traditional wireline network that supports the wireless deployment is fiber. But that is about to change. DOCSIS technology is one of the key enablers for mobile deployments because of its near-ubiquitous availability in urban and suburban areas. Residential femto, strand-mounted small cells, as well as SMB inside-out will provide inside and outside coverage economically compared to fiber.

Since the cable industry first specified the DOCSIS standard in 1997, the technology has evolved through five generations of progressive improvements over several key performance criteria, including capacity and latency. Table 7 shows the capabilities of the most recent DOCSIS standards. The currently deployed DOCSIS 3.1 is capable of supporting multi-gigabit per second downstream speeds. As MSOs move to reclaim spectrum previously used for traditional video services, upstream spectrum can be significantly increased by moving from the low split of 42 MHz to the high split of 204 MHz. Multi-gigabit per second of speeds on the upstream can be reached this way in the near term. In the longer horizon, DOCSIS 4.0, currently being specified, is expected to provide greater upstream speeds.

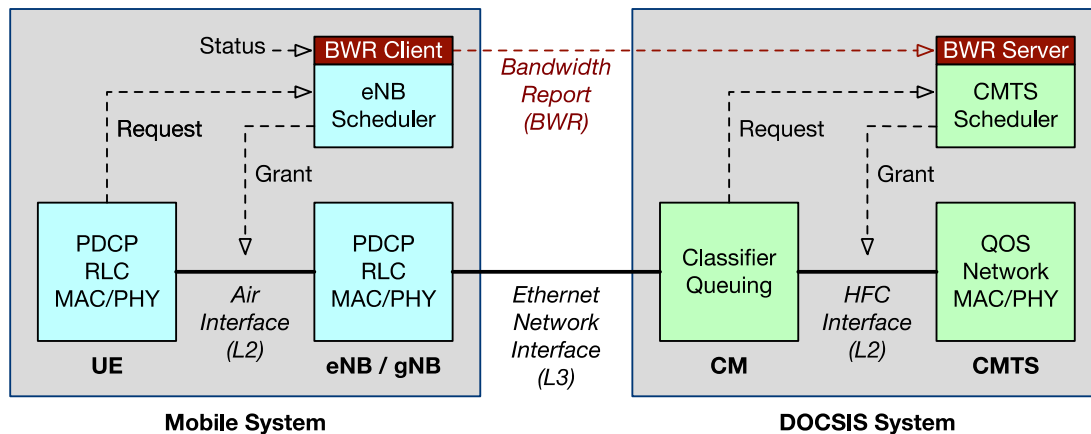
**Table 7 – DOCSIS Capabilities**

Requirements	D3.1 today (2020)	D3.1 max (future)	D4.0 (2023-2024)
	Shared spectrum with video	Full spectrum through video reclamation	Extending to 1.8 GHz, possibly 3 GHz
Downstream spectrum	54 - 1002 MHz	258 - 1218 MHz	602 - 1794 MHz
Upstream spectrum	5 - 42 MHz	5 - 204 MHz	5 - 492 MHz
DS Capacity	8.5 Gbps	8.6 Gbps	10.8 Gbps
US Capacity	0.1 Gbps	1.4 Gbps	3.7 Gbps
Latency	Best Effort : 5 - 50 ms.	With LLX/ CTI : 1 - 2 ms (can be further reduced)	
Synchronization	Frequency sync only No time sync	Frequency + time sync through DTP	

### 5.2. Low Latency Xhaul (LLX)

DOCSIS technology supports a variety of scheduling mechanisms to meet the latency needs. The most commonly used best effort scheduler can deliver a typical latency of around 10 to 15 milliseconds (ms), but is dependent on the channel loading condition (see Table 7). DOCSIS also natively supports real-time polling (RTPS) and proactive grant service (PGS). These schedulers intend to reduce the request-grant delay that is typical in any point-to-multipoint scheduled systems. But they have some drawbacks, such as not able to achieve low enough latency needed for mobile xhaul, or incurring too much bandwidth overhead.

To address these drawbacks and to better support mobile xhaul over the DOCSIS network, Cisco and CableLabs co-invented a mechanism that pipelines the scheduler operations of the mobile and the DOCSIS systems. The pipelining mechanism forms the basis of the Low Latency Xhaul (LLX) technology, which has been standardized by CableLabs [6].



**Figure 56 – Scheduler Pipelining with Bandwidth Report (BWR)**

In a nutshell, LLX uses the decisions made by the mobile scheduler to inform the CMTS scheduler what is about to happen. By doing so, LLX creates a low latency transport for mobile traffic.

As shown in Figure 56, mobile and DOCSIS systems are both point-to-multipoint scheduled systems. This means both systems have an inherent latency due to request-grant delay in the upstream. In LLX, that latency is incurred once in the mobile system. The results of the request-grant process, in the form of a BWR message, are then passed to the DOCSIS system so the CMTS can grant the CM directly without waiting for a native layer 2 DOCSIS request.

Let's look at an example.

1. The UE has an application that wants to send 1000 bytes. It sends a request to the eNB scheduler.
2. The eNB scheduler responds and says that the UE may send the 1000 bytes 8 ms from a reference time.
  - a. The eNB scheduler, now that it knows what is about to transpire on its air interface, makes a determination of what will happen across the network interface that it shares with the DOCSIS system. In our example, the eNB adds 1 ms of engineering margin to cover any buffering and internal path delays.
3. The eNB sends a BWR message to the CMTS system that says that 1000 bytes will be arriving on the shared network port 9 ms from the reference time.
  - a. The CMTS scheduler, now that it knows when the bytes will arrive in the CM, determines when it wants to send a grant to that CM. In this example, it adds 1 ms of engineering margin to cover any buffering or scheduling jitter.
4. The CMTS sends a DOCSIS MAP to the CM at the correct time telling the CM to transmit the 1000 bytes 10 ms from the reference time.

The net result is that the latency of the DOCSIS system is effectively reduced by hiding it under the mobile system. In theory, LLX should be able to achieve near-zero latency on the DOCSIS upstream. In practice, one to two milliseconds of engineering margin is added.

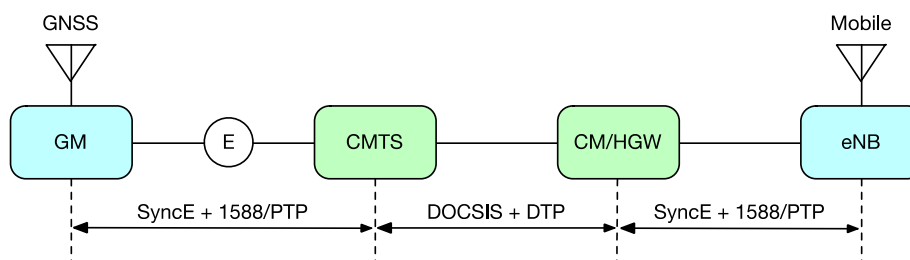
Numerous test results have been previously published using a physical testbed and reported that BWR achieves one to two milliseconds of DOCSIS upstream latency in a variety of channel loading conditions up to 70% on the DOCSIS network. Details of how LLX works and lab trial results can be found in [11][12][13][14][15].

### 5.3. Synchronization and Timing

Depending on the type of mobile deployment, frequency only, or frequency and phase, synchronization is required for the small cell. Timing and synchronization requirements for FDD, TDD, LTE, and 5G are shown in Table 4.

Table 5 in Section 4.1.4.4 lists common options to consider when it comes to support timing and synchronization. Network-supported timing using Precision Time Protocol (PTP) is needed for indoor deployments and to lower the cost of the small cells.

The DOCSIS network is asymmetrical. If the PTP protocol is sent over-the-top of the DOCSIS network, it may incur variable buffer delay which can cause packet delay variation (PDV) and large time transfer errors. For LTE FDD, PTP over-the-top may be a workable solution with some mitigation work, such as assigning PTP packets with higher priority DiffServ code point (DSCP). But the mitigation may not be enough for LTE and 5G TDD.



**Figure 57 – 1588 Timing over DOCSIS Network with DTP**

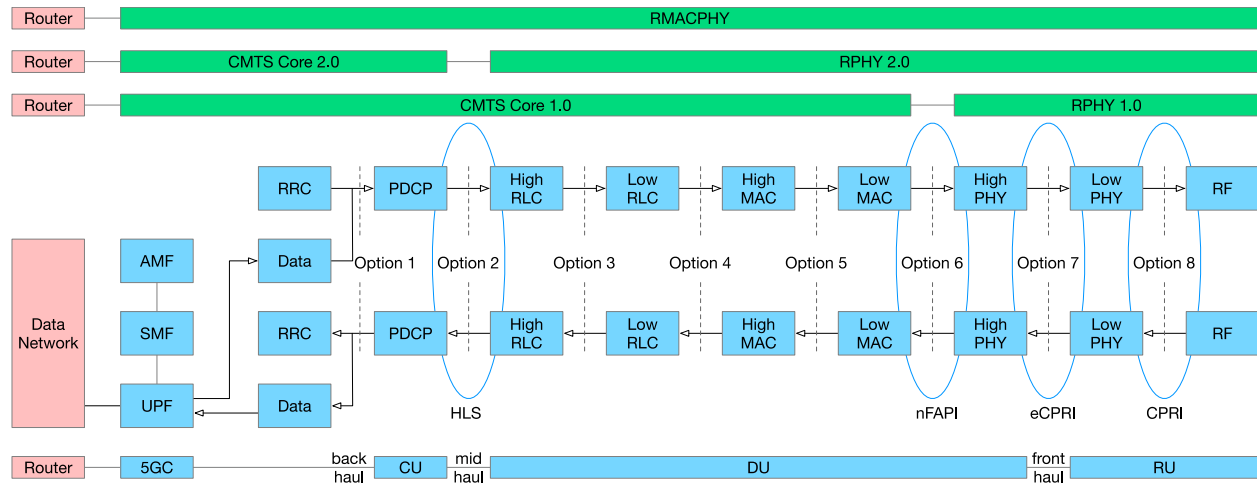
To meet the TDD requirement, a better way to carry PTP over the DOCSIS network is through the DOCSIS Time Protocol (DTP), as shown in Figure 57. In a nutshell, timing from the global navigation satellite system (GNSS) is received by a primary reference time clock (PRTC) that generates the PTP messages. The PTP messages are sent through one or more Ethernet switches that support PTP. The PTP timing domain is terminated by the CMTS.

The DOCSIS system is already a synchronous network with its own timestamp. The CMTS's job is to align the DOCSIS timestamp with the PTP timestamp. The DTP algorithm is run between the CMTS and the CM to compute the one-way downstream delay. Upon receiving the DOCSIS timestamp as part of the normal DOCSIS operations. The CM adds the one-way delay to the DOCSIS timestamp. The CM passes on the recomputed timestamp and appears as a PTP master to the small cell downstream.

Details of the DTP algorithm and preliminary test results can be found in [16][17][18]. DTP is now part of the synchronization specification standardized by CableLabs [19].

## 5.4. DOCSIS DAA for Mobile People

This section briefly discusses the splits for the 5G radio access network (RAN) and then compares them to the choices made by DOCSIS when it split its access network with the distributed access architecture (DAA). The contrast of the two architectures are shown in Figure 58.



**Figure 58 – Mobile RAN Splits with DOCSIS DAA**

The 3GPP task force defined eight splits in the RAN architecture and then picked higher layer split (HLS), while the O-RAN Alliance further specified the interoperability required to enable lower layer split (LLS). The original common public radio interface (CPRI) was a raw digital to analog conversion style of interface. It has a very large bit rate and is considered semi-proprietary. The enhanced CPRI (eCPRI) lowered the bandwidth to make it fit better on 10 Gbps fiber links and was adopted as the LLS / option 7 split between the radio unit (RU) and the distributed unit (DU). This interface is also referred to as the fronthaul interface.

The DU contains layer 2 framing and the uplink scheduler. The DU is intended to be located near the RU location. Typically, the RU is outdoors and the DU is indoors. A DU product may service multiple RU products. The network side of the DU is a packet interface with a data throughput slightly above the payload rate due to only having payload encapsulation and signaling. The split between the DU and the centralized unit (CU) is done with option 2 and is referred to as the HLS.

The first architecture for DAA was Remote PHY (RPHY) [20][21]. Remote PHY is well defined as an open standard [22]. It is a shipping product with multiple manufacturers and operators deploying RPHY with several million attached devices. The goal of RPHY is to centralize software and distribute the radio frequency (RF) hardware. Because the DOCSIS protocol has encryption built into it, it was convenient to put the DOCSIS layer 2 framer centrally to ensure the backhaul link was encrypted. This would be the equivalent of 3GPP's option 6. In fact, before option 2 and option 7 were chosen for the mobile RAN, there was a proposal from Cisco for an option 6 network functional application platform interface (nFAPI) in the Small Cell Forum (SCF) [23].

After Remote PHY was defined, and broader architecture was defined called the flexible MAC architecture (FMA). The first phase of FMA is the remote MAC and PHY (RMACPHY) which is really a full layer 2 cable modem transport system (CMTS) that includes the entire CMTS Core, or the equivalent of a full EPC,



located in the remote fiber node. In the mobile world, FMA resembles incorporating the 5G Core into an integrated gNB, both located the remote node. At the time of writing, FMA is a pre-standard draft.

There is a next generation of Remote PHY, informally called RPHY 2.0 [24], that is proposed but not yet standardized. A RPHY 2.0 Device (RPD 2.0) provides DU-like capabilities that include the DOCSIS upstream scheduler [25] and the DOCSIS framer. The first version of RPHY was designed about a remote RF device and a centralized physical set of DOCSIS and video cores that were located within a 100 mile (160 km) radius of the RPD. By contrast, the second version of RPHY will be designed to put latency sensitive signaling such as the upstream scheduler into the remote node. Doing so allows distances in excess of 100 miles to cloud-based cores that are located at the service provider (SP) edge or in true cloud such as Amazon or Azure.

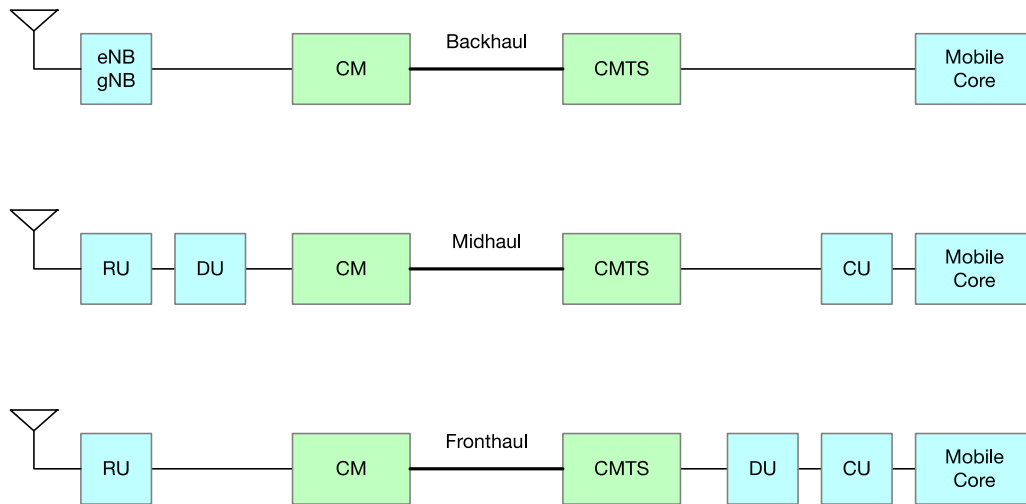
## **5.5. DOCSIS for Mobile Xhaul**

Similar to the DOCSIS protocol, LTE and 5G have their own protocol stack. LTE eNB or 5G gNB are full-stack integrated units. In 5G, the RAN can be further functionally decomposed into a central unit (CU), a distributed unit (DU), and a radio unit (RU). Mobile standards organization 3GPP specified the option 2 split, which splits the PDCP and above layers into the CU while the layers below remain in the DU and RU. The ORAN Alliance defined option 7-x split, in which the RU consists of the RF and a portion of the PHY, while the layers above remain in the DU and the CU. A backhaul, midhaul, or fronthaul network, collectively known as xhaul, interconnects the different functional components together.

Functionally, the DOCSIS network can interconnect the small cells as shown in Figure 59. However, the requirement on capacity and latency vary between backhaul, midhaul, and fronthaul. Backhaul and midhaul are both based on an IP encapsulation of the original mobile transport. Thus, the bandwidth requirement on the transport network roughly matches the mobile traffic rate.

The latency requirement for backhaul is based on the application. Additional service-level agreement (SLA) can be specified by the mobile operator. Midhaul latency is less deterministic. Some standards have defined it to be less than 10 ms [26], while some vendors require one to three milliseconds of CU-DU latency. LLX can be implemented to ensure these requirements can be met, as well as providing better latency performance, particularly for latency-sensitive flows.





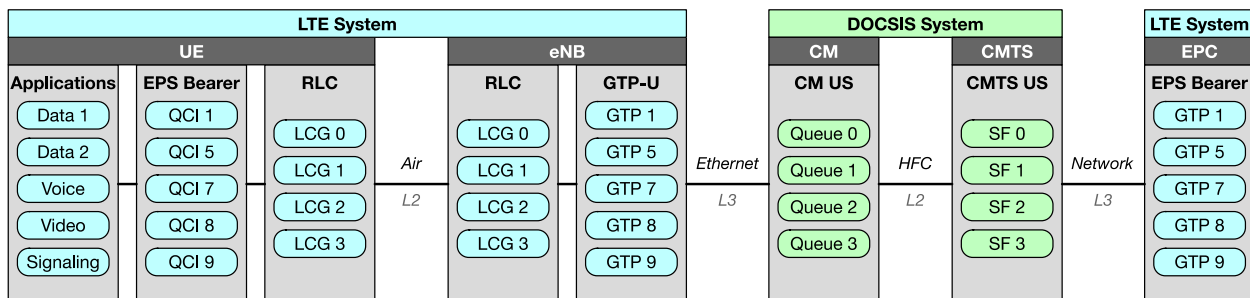
**Figure 59 – Mobile Xhaul over DOCSIS Architecture**

Fronthaul is much more difficult to support over the DOCSIS network. Studies have shown that the eCPRI-based fronthaul transport needs significantly more bandwidth, overhead, and one-way latency in the neighborhood of 250 microseconds between the RU and DU [20]. Even with LLX, it will be difficult to reduce the DOCSIS latency to this level. Because of the stringent requirements, fronthaul is better carried over fiber.

## 5.6. Common Quality of Service (QoS) Framework

The DOCSIS network is a finite pipe with multiple endpoints sharing resources. Rather than dedicating resources to meet the peak capacity all the time, the DOCSIS network is designed to meet the bandwidth demand most of the time. Traffic is separated into multiple flows. During times of congestion, latency-sensitive flows such as signaling or 5G ultra-reliable low-latency communication (URLLC) traffic are sorted into separate queues and are generally served before latency-tolerant flows. This is the goal of quality of service.

The mobile system is also a point-to-multipoint system where resources are shared among users. It has its own set of QoS rules and queue configurations that may be different from the DOCSIS network.



**Figure 60 – Common Quality of Service Framework for MBH over DOCSIS**

To ensure consistent treatment of traffic when they move across the mobile and DOCSIS networks, a common QoS framework as shown in Figure 60 between the mobile system and the transport system needs to be supported. There are many variations of how this could be done, but fundamentally, it should:


1. Use the same number of queues in the transport system as there is in the mobile system
2. Use the same classifier mechanism in the transport system as there is in the mobile system
3. Use the same policy/queue-weighting mechanism in the transport system as there is in the mobile system

Details of the common QoS framework can be found in [6][9][27].

## 5.7. A 5GC View of Convergence

The advent of 5G promises several key enhancements to previous mobile technology generations including faster speeds, lower latency, and increased service velocity which enables new use cases such as enhanced mobile broadband, massive IoT and mission critical enterprise applications. Along with new enhancements in the RAN, the 5G core (5GC) has been rearchitected to support these new use cases.

The 5GC architecture includes service-based design concepts, on-demand network slicing and service orchestration, and cloud native design principles such as web-based control plane protocols, microservices and container orchestration. One of the most innovative architectural concepts is wireless and wireline convergence (WWC); allowing different access networks such as Wi-Fi, cable, and fixed broadband networks to interwork with the 5GC as shown in Figure 61. Achieving 5G convergence enables new use cases and consistent subscriber quality of experience (QoE) while reducing cost of ownership.



**Figure 61 – 5G Convergence**

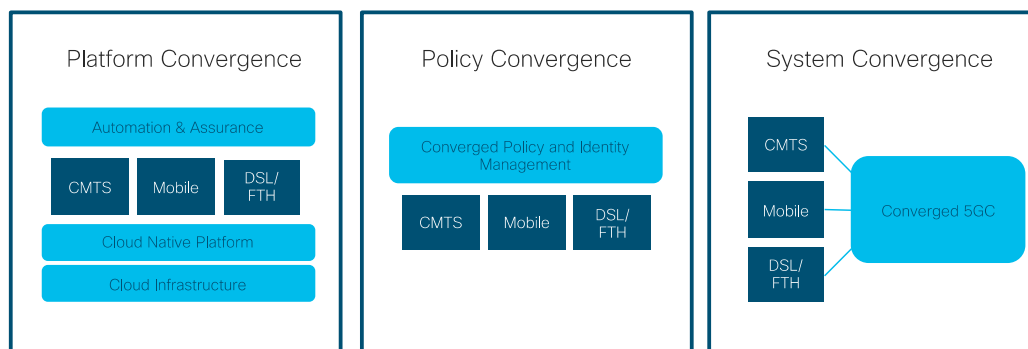
5G convergence enables multi-access operators to harmonize subscriber services across access networks. As subscribers move between their home broadband and mobile networks, a converged core enables an improved quality and consistency of user experience. This includes use cases such as call continuity so that voice and video calls are seamlessly handed off between connections. Policy based services, such as bandwidth speed and latency tiers, as well as parental controls, security and content filtering can also be consistently enforced.

Convergence also enables connection redundancy which has become increasingly important for enterprises and consumers as users are working and schooling from home during the pandemic. For example, a 5G fixed wireless access service can be coupled with a fixed broadband service to enable service resiliency in an active/active or active/standby design.

Arguably as important as QoE, convergence allows operators to reduce their capital and operating expenditures. Historically, each access network has been deployed and operated as a silo in the SP network

with different infrastructure, networking and applications. Network convergence allows SPs to collapse these silos into a common, converged network where the edge and core infrastructure, applications and orchestration systems are common across access networks. Convergence also results in OpEx reductions as existing siloed networks collapse into a common converged core.

In addition to the network, convergence allows the consolidation across all aspects of the SP operations extending into OSS/BSS systems (not covered in this paper). There are three common approaches to converging subscriber QoE across access technologies. The platform, policy, and system level convergence approaches represent different steps towards the path to convergence. Each of these steps, along with their benefits, are shown in Figure 62 and are described in the following paragraphs.



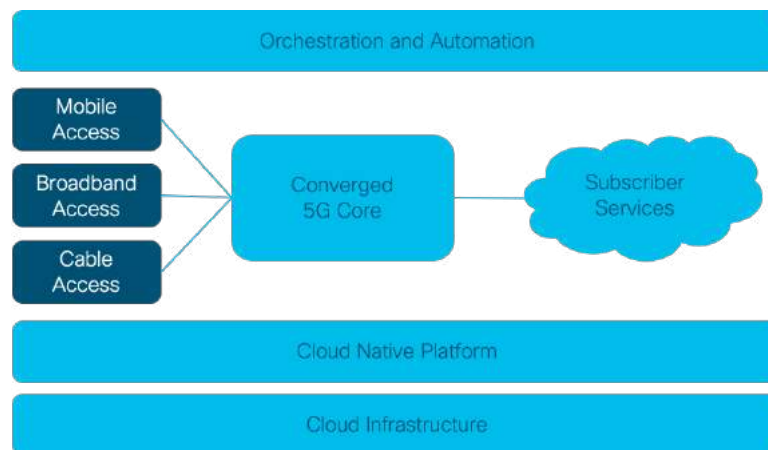
**Figure 62 – Platform, Policy, and System Convergence**

Platform convergence utilizes a common platform to achieve operational and software development efficiencies. The converged platform includes both physical and virtual infrastructure, as well as cloud native application technologies to deploy and operate the different access systems. It provides common compute and virtualization technologies, such as microservices, containerization, Kubernetes, and cloud assurance and automation systems, to provide a consistent operations experience for SP operation teams. Platform level convergence does not require each access function to be converged by a common system and protocol architecture, but means common tooling is used for deploying, automating, orchestrating and assuring each access function.

Policy level convergence enables common policy and identity control plane systems but does not try to converge the network user plane components. This approach enables policy coordination across access networks where each access network is responsible for enforcing policies provided by the common policy layer. Therefore, although there are separate transport and user planes for each access network, the subscriber experience is maintained by the converged policy layer across networks, e.g., a parental control service can be enforced on both a broadband and mobile network. Many SPs are already on the path to subscriber and policy level convergence as a first step towards transport and user plane convergence across accesses.

System level convergence of the end to end architecture is not a new concept and has been a topic of previous technology generations – but as an afterthought. Convergence was included as part of the 5G standalone design from its inception and allows existing access networks to be integrated with a common converged core based on an interworking approach. This allows existing access technologies to be integrated into the converged core as opposed to requiring each access to be reimplemented. WWC is being designed in standards as a cooperation across 3GPP, BBF, CableLabs and other standards bodies. WWC is the long-term strategy for most multi-access providers, but it comes with significant cost and complexity for brownfield networks. Many MSOs are deploying greenfield 5G networks and are looking to provide

convergence with their Wi-Fi and DOCSIS networks over time. Figure 5 in the earlier section provides a detailed view of the converged access architecture proposed for combining mobile and cable access.



**Figure 63 – Common Cloud Native Platform**

To successfully navigate the migration to a converged core, operators must determine which convergence type can be justified based on their business. Additionally, operators must structure their product and operations teams to align with their chosen level of convergence. Ultimately, we expect the industry to pursue a common cloud native platform architecture as shown in Figure 63 to simplify and automate operations, common policy and subscriber components to unify the subscriber QoE, and a fully converged core to harmonize the services and applications into a common network architecture.

## 5.8. Managing the RAN with YANG

One of the key principles of the vRAN architecture is to establish open, standard interfaces, facilitating a transition from today's single vendor, monolithic RAN solutions, to a competitive, multi-vendor environment where functionality can be sourced from a variety of vendors that address the unique needs of MSO deployments around the world. While control and user-plane protocols are naturally required to be interoperable, an often-overlooked aspect is the integration "tax" required to deliver a fully orchestrated multi-vendor system.

Crucially, there are no procedures in 3GPP to ensure that all the parameters necessary for 5G can be configured using the 3GPP management specifications. This gap is being filled by the Open RAN (O-RAN) alliance that has already defined the use of native YANG models for configuring its distributed Radio Unit, the so-called O-RAN RU (O-RU). Not only does the use of native YANG models ensure the easiest route to full multi-vendor interoperability, but it also eases the integration of the management for the lower layer split deployment into existing systems. This has been demonstrated by multiple multi-vendor deployments of the O-RAN fronthaul interface.

YANG (RFC 7950) is a modelling language that was initially adopted by the xRAN Forum in March 2018, to model the configuration and operational state of its 5G Radio Unit, together with defining remote procedure calls (RPCs) for supporting tasks like software management, and notifications for indicating xRAN defined alarms. In 3GPP Release 16, YANG has been adopted as a new protocol-specific solution-set that leverage a protocol-neutral Network Resource Model (NRM). Now incorporated into the O-RAN Alliance, the O-RU specifications use YANG to define syntax, relationships and constraints between the data, enabling operators of O-RAN's open fronthaul to validate configuration data against the model before committing the configuration of the O-RAN Radio Units.

Recognizing that O-RAN Radio Unit suppliers need to be able to support vendor differentiation, the YANG models are extensible, allowing them to be augmented to support enhanced vendor-specific functionality, while simultaneously ensuring baseline multi-vendor interoperability of the standardized functionality defined by O-RAN.

The use of augmented IETF standard YANG models, together with O-RAN specific models, lays the foundation for cross-domain orchestration of the RAN with other domains that have already adopted NETCONF/YANG. Importantly, the definition of the transport interfaces in O-RAN leverages the IETF standard defined YANG models which should then facilitate the use of common tooling to be used across transport and RAN domains.

These same YANG models also provide the foundation for model-driven telemetry. Instead of ill-documented command line interface (CLI) or poor scalability of simple network management protocol (SNMP), model-driven telemetry enables data to be streamed from network devices continuously using a push model and provides near real-time access to operational statistics. Applications can subscribe to specific data items they need, by using standard-based YANG data models over NETCONF-YANG.

## 6. Convergence – A Vision of What is to Come

### 6.1. Integrated and Converged HFC, DOCSIS and Mobile Network

To converge or not; the answers may be in the business case. The operators have the choices of levels of convergences that they are comfortable with, starting with spending little CapEx and loosely coupling the two systems, to spending more capex and tightly coupling the two systems.

Up to this point in the paper, we have looked at a broad framework for convergence, the visions of convergence from multiple operators, and specific technologies of convergence. In this section, we will pull the basics of transport and infrastructure convergence into one vision.

Figure 64 represents a vision of the points of convergence that might ultimately take place:

- transport network convergence – common CIN and coax carry mobile, DOCSIS, and PON traffic
- common cloud platform – run apps on common edge
- converged cloud native functions (CNF) – common policies, common user plane functions

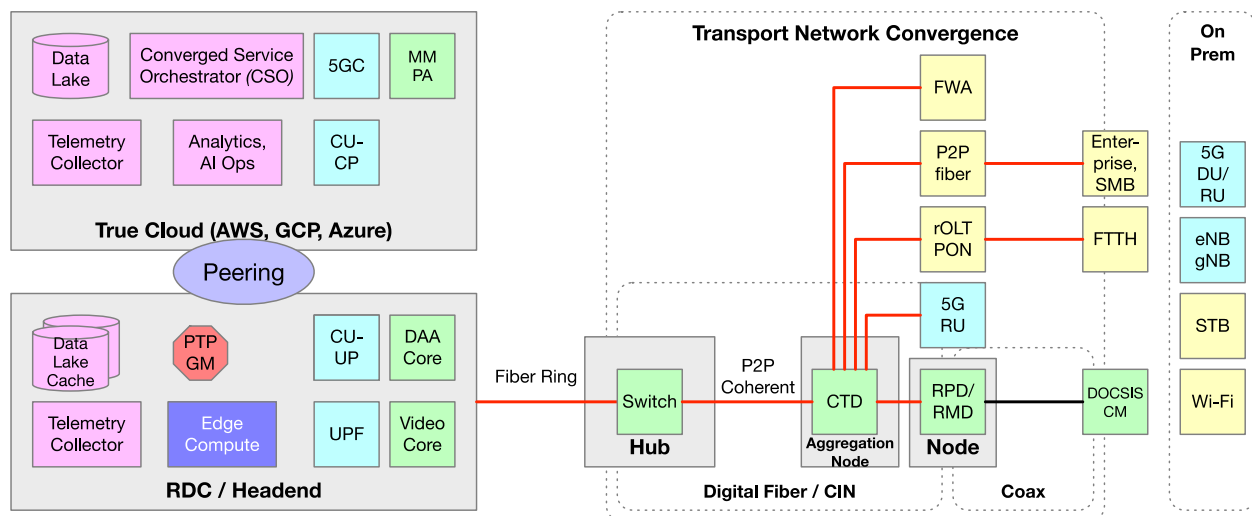


Figure 64 – DOCSIS-Mobile Network Convergence

### 6.2. Converged Transport of Mobile Xhaul over DOCSIS

Mobile xhaul over DOCSIS is a form of transport network convergence that does not require significant CapEx spending by allowing the MSOs to reuse their HFC plant to carry mobile traffic. All MSOs deploy integrated CMTSs (I-CMTS) today. The DOCSIS network today can be used to carry mobile traffic. Some of the technologies described in a section above such as LLX and DTP can enable enhanced and optimized mobile xhaul deployments.

### **6.3. Converged Transport with Common CIN**

The next level of convergence takes place as MSOs modernize their DOCSIS network to meet the capacity demands on the HFC network.

One of the classic tools in the MSO toolbox to increase plant capacity is node segmentation. When an operator decides to segment a node, there is an opportunity to transform the cable-specific analog optics technology into digital. This is one of the benefits that the remote PHY (R-PHY) architecture brings. By disaggregating the PHY and RF-generating components of the DOCSIS stack to the node location, operators can replace their old analog lasers with new digital lasers, and in the process, increase capacity and open up the optical network. This upgrade also means that a generic Ethernet switch can be deployed in the aggregation node, connecting not just the remote PHY device (RPD) or remote MACPHY device (RMD), but also small cells and PON networks.

The analog to digital optics transformation of the HFC plant requires the buildout of the converged interconnect network (CIN), which includes the layer 2 switches that enable the same fiber network to be used by DOCSIS, mobile, and PON. This is another important aspect of transport network convergence.

### **6.4. Common Cloud Platform**

As mobile and cable vendors adopt cloud native technologies, many software processes that run on dedicated hardware platforms today can be run as cloud native applications on generic servers. This can include upper layers of DOCSIS, 5G core, and some of the RAN functions. These processes from the DOCSIS network and the mobile network can run on the same server complex.

With control and user plane separation (CUPS) architecture, user plane processing that requires lower latency can be split from the more latency-tolerant control plane functions. These data plane functions, including CU-user plane (CU-UP), 5G's user plane function (UPF), and DOCSIS data plane (DP), can be run as software process at the edge, whereas CU-CP, much of the 5GC, and DOCSIS CP can be relocated to the cloud for large scalability.

To operate and monitor a network, telemetry data need to be collected and processed as part of the service assurance framework. A converged data lake and telemetry collector can be used for both DOCSIS and mobile networks.

Traditionally, MSOs own and operate a variety of real estate besides the fiber and coax runs. Of these resources, headends or regional data centers (RDC) can most likely be scaled to run data center applications. It is here that DOCSIS and mobile software processes can be located.

### **6.5. Virtual and Cloud Native Functions**

When software was first moved to server complexes, which was primarily with mobile architectures, it was done using virtualized network function (VNF). In virtualization, existing code can be ported from a physical platform, such as a physical EPC, into a virtual machine (VM) on a server. The advantage of this approach was preservation of code and time to market. The goal was cost reduction and a use of generic servers. Unfortunately, the costs savings did not materialize, and the code was still old code with maintenance, complexity, and scalability challenges.

The new way of writing code involves cloud native functions (CNF). In cloud technologies, old code needs to be rearchitected and rewritten, often using newer languages such as the Go language. The code is often written to be stateless so that processes can quickly restart when the crash without the loss of a session.

The code is partitioned into microservices which are then placed into Docker containers. Kubernetes (K8S) is used to do workload placement of these containers into a server-based system. The result is a system that is highly elastic and highly resilient. This represents new functionality which has tremendous market value and makes the move to servers based on an increase in value, rather than just cost reduction.



## 7. Conclusion

Convergence is equally important now as providing a clear path to the end of the tunnel. On the road to that end deployment state, tools controlling the operating costs are evolving today, including the movement to cloud-native architecture. In addition to virtualization, the following are the common technologies to enable mobile deployment identified by most MSOs:

- optimizations to enable mobile xhaul over DOCSIS such as DTP, LLX, and orchestration
- the upcoming HFC buildout from analog optics to a digital fiber network
- the CIN, whether be it simpler DWDM multiplexer or Ethernet switch for better scalability
- 5G NSA core for MSOs operating existing LTE networks, while transitioning to 5G NR
- DSS for migration from LTE RAN to 5G NR in existing LTE spectrum
- DSDS to support better MVNO economics with better data offloading
- CUPS to enable virtualization while still providing lower user-plane latency
- tools for service convergence: common policy, common subscriber management
- model-driven telemetry and AI operations for network automation

*Today's cable operators are tomorrow's mobile operators.*

Subscribers, operators, and networks are converging. Ignoring that reality is no longer an option. There is now a whole new world of mobile operators, just in time for 5G.

*Every great wireless network needs a great wireline network.*

Of all the variations amongst the operators, there are some commonalities. Most certainly, the end state for network deployment will not be a single physical fiber-based network to rule them all, but a combination of physical networks that involve fiber, coax, and wireless, to support DAA, FTTH, FWA, mobile deployments.

Everything reduces to an IP network with different edge connectivity with common services, management and provisioning.

Remember the telephone network. It connected voice end points called telephones. The telephone network does not significantly exist anymore. The same will be true for cable and mobile networks. They were separate networks because they were owned and run separately, and they may have had unique backbone and edge requirements. But that is or has already changed.

Today, the mobile network has really become an IP over Ethernet network that connect radio gateways to specific user plane points and run by control and management applications running on generic servers. The mobile network is becoming a software-based design with an IP transport terminated in radios for end-point connectivity.

The same transformation is happening to the cable network. Once a closed HFC plant, with DAA there will be IP over Ethernet to the neighborhood with the RPD acting as a radio attachment point that drives the last mile of coax. Even DOCSIS is just a form of Ethernet over coax and acts like a fiber extension in DAA. That is why either fiber or coax can be used for mobile xhaul. Mobile xhaul over DOCSIS provides MSOs the economic advantage for deploying wireless networks. It is an obvious convergence opportunity.

Careful decisions have to be made on where convergence adds value and simplicity, or where it adds cost and complexity. Which decisions make money and which ones lose money. It is often elegance versus execution. But that is just the challenge of getting innovation right.

So, in the most simplistic of terms, both cable and mobile are really composed of different radios on a common IP network trying to do the same thing, and that is to deliver a common set of services to a subscriber no matter how they are connected. The job of convergence is to make this reality come true with both a common service model and an efficient and economical network and application infrastructure.

Let's go make it happen.

# Abbreviations

<b>2T2R</b>	2 transmitters 2 receivers
<b>3GPP</b>	third generation partnership project
<b>5GC</b>	5G core
<b>ADSL</b>	asynchronous DSL
<b>AFC</b>	automated frequency control
<b>AMAP</b>	Africa, Middle East, Asia-Pacific
<b>ANCTD</b>	aggregation node with coherent termination device
<b>AOI</b>	area of interest
<b>AP</b>	access point
<b>AR</b>	augmented reality
<b>ARPU</b>	average revenue per user
<b>ATSSS</b>	access traffic steering, switching, and splitting
<b>BNG</b>	broadband network gateway
<b>BSS</b>	business support systems
<b>BWR</b>	bandwidth report
<b>C-RAN</b>	centralized radio access network
<b>CCAP</b>	converged cable access platform
<b>CBRS</b>	citizen broadband radio service
<b>CBSD</b>	CBRS device
<b>CIN</b>	converged interconnect network
<b>CLI</b>	command line interface
<b>CM</b>	cable modem
<b>CMTS</b>	cable modem termination system
<b>CNF</b>	cloud native function
<b>CO</b>	central office
<b>CP</b>	control plane
<b>CPE</b>	customer premise equipment
<b>CPRI</b>	common public radio interface
<b>CSP</b>	communication service provider
<b>CTD</b>	coherent termination device
<b>CU</b>	central unit
<b>CUPS</b>	control and user plane separation
<b>DAA</b>	distributed access architecture
<b>DC</b>	data center
<b>DFS</b>	dynamic frequency selection
<b>DL</b>	downlink
<b>DOCSIS</b>	data over cable system interface specification
<b>DP</b>	data plane
<b>DS</b>	downstream
<b>DSCP</b>	differentiated services code point
<b>DSDS</b>	dual SIM dual standby
<b>DSL</b>	digital subscriber line
<b>DSS</b>	dynamic spectrum sharing
<b>DTP</b>	DOCSIS Time Protocol
<b>DU</b>	distributed unit
<b>DWDM</b>	dense wave division multiplexing

<b>eCPRI</b>	enhanced CPRI
<b>eMBB</b>	enhanced mobile broadband
<b>eNB</b>	eNodeB
<b>EPC</b>	evolved packet core
<b>eSIM</b>	embedded SIM
<b>FCC</b>	Federal Communications Commission
<b>FMA</b>	flexible MAC architecture
<b>FMC</b>	fixed mobile convergence
<b>FO</b>	fiber optical, fiber optics
<b>FTTC</b>	fiber to the cabinet
<b>FTTH</b>	fiber to the home
<b>FWA</b>	fixed wireless access
<b>gNB</b>	gNodeB
<b>GNSS</b>	global navigation satellite system
<b>GPS</b>	global positioning system
<b>HeBGW</b>	home eNodeB gateway
<b>HeMS</b>	home eNodeB management
<b>HeNB</b>	home eNB
<b>HFC</b>	hybrid fiber-coaxial
<b>HHP</b>	households passed
<b>HLS</b>	higher layer split
<b>HMNO</b>	hybrid mobile network operator
<b>ISED</b>	department of Innovation, Science and Economic Development
<b>IoT</b>	Internet of things
<b>IPBB</b>	IP backbone
<b>iWinS</b>	intelligent wireless network steering
<b>LAN</b>	local area network
<b>LLS</b>	lower layer split
<b>LLX</b>	low latency xhaul
<b>LTE</b>	long term evolution
<b>MDU</b>	multi-dwelling units
<b>MIMO</b>	multiple in multiple out
<b>MNO</b>	mobile network operator
<b>ms</b>	millisecond
<b>MSO</b>	multiple system operator
<b>MVNO</b>	mobile virtual network operator
<b>NB-IoT</b>	narrowband IoT
<b>nFAPI</b>	network functional application platform interface
<b>NFV</b>	network function virtualization
<b>NGA</b>	next generation access
<b>NPRM</b>	notice of proposed rulemaking
<b>NR</b>	new radio
<b>NRM</b>	network resource model
<b>NSA</b>	non-standalone
<b>NTP</b>	network time protocol
<b>ODN</b>	optical distribution network
<b>OEM</b>	original equipment manufacturers
<b>OLT</b>	optical line termination
<b>ONT</b>	optical network termination
<b>opex</b>	operating expense

<b>O-RAN</b>	Open RAN
<b>O-RU</b>	O-RAN RU
<b>OS</b>	operating system
<b>OTT</b>	over the top
<b>PON</b>	passive optical network
<b>PDV</b>	packet delay variation
<b>PGS</b>	proactive grant service
<b>PLMN</b>	public land mobile network
<b>PRTC</b>	primary reference time clock
<b>PTP</b>	precision time protocol
<b>QoE</b>	quality of experience
<b>QoS</b>	quality of service
<b>RAN</b>	radio access network
<b>RDC</b>	regional data center
<b>RF</b>	radio frequency
<b>RG</b>	residential gateway
<b>RGW</b>	residential gateway
<b>RMACPHY</b>	remote MAC and PHY
<b>RMD</b>	RMACPHY device
<b>rOLT</b>	remote optical line termination
<b>RPD</b>	remote PHY device
<b>RPHY</b>	remote PHY
<b>RTPS</b>	real-time polling service
<b>RTT</b>	round trip time
<b>RU</b>	radio unit
<b>SA</b>	standalone
<b>SCF</b>	Small Cell Forum
<b>SDN</b>	software defined network
<b>SeGW</b>	security gateway
<b>SFP</b>	small form-factor pluggable
<b>SG</b>	service group
<b>SHDSL</b>	single-pair high-speed DSL
<b>SLA</b>	service level agreement
<b>SMB</b>	small and medium business
<b>SNMP</b>	simple network management protocol
<b>SOHO</b>	small office home office
<b>SON</b>	self-optimizing network
<b>SP</b>	service provider
<b>TCO</b>	total cost of ownership
<b>TDD</b>	time division duplexing
<b>U-NII</b>	unlicensed national information infrastructure
<b>UE</b>	user equipment
<b>UL</b>	uplink
<b>UP</b>	user plane
<b>UPF</b>	user plane function
<b>US</b>	upstream
<b>USP</b>	user services platform
<b>URLLC</b>	ultra-reliable and low-latency communications
<b>vCCAP</b>	virtualized CCAP
<b>VDSL</b>	very high-speed DSL

<b>vEPC</b>	virtualized EPC
<b>VM</b>	virtual machine
<b>VNF</b>	virtual network function
<b>VoWiFi</b>	voice over Wi-Fi
<b>VR</b>	virtual reality
<b>vRAN</b>	virtualized RAN
<b>WDM</b>	wavelength division multiplexing
<b>WWC</b>	wireless and wireline convergence

## Bibliography & References

- [1] MoffettNathanson, “Wireless Q2 2020: Chasing T-Mobile”, August 2020.
- [2] Jennifer Andreoli-Fang, “DOCSIS-Mobile Convergence,” Broadband Library, summer edition 2019. [\[link\]](#)
- [3] “Non-roaming architecture for UE behind 5G-RG using trusted N3GPP access”, 3GPP TS 23.316 V16.3.0 (2020-03) Figure 4.10-1
- [4] Jennifer Andreoli-Fang, Alon Bernstein et al., “Network Convergence”, *SCTE Fall Technical Forum*, Oct, 2018. [\[link\]](#)
- [5] John T. Chapman, “Small Cell Traffic Engineering,” *SCTE Fall Technical Forum*, Oct, 2020. [\[link\]](#)
- [6] “Low Latency Mobile Xhaul over DOCSIS Technology,” CM-SP-LLX, CableLabs. [\[link\]](#)
- [7] Dave Morley, “5G Small Cells and Cable – Realizing the Opportunity,” *SCTE Fall Technical Forum*, October, 2018. [\[link\]](#)
- [8] Damian Poltz, “HFC and Wireless – Cable’s Convergence Advantage”, *CableLabs Summer Conference*, Aug, 2019
- [9] Cisco Annual Internet Report (2018–2023) White Paper, Cisco. [\[link\]](#)
- [10] “5G Implementation Guidelines: NSA Option 3,” 28 March 2019, GSMA.
- [11] Jennifer Andreoli-Fang, John T Chapman, Tong Liu, Damian Poltz, “Blueprint for Mobile Xhaul over DOCSIS,” *SCTE Fall Technical Forum*, Sep, 2019. [\[link\]](#)
- [12] John T. Chapman, Jennifer Andreoli-Fang, “Low Latency Techniques for Mobile Backhaul over DOCSIS,” *SCTE Fall Technical Forum*, October, 2017. [\[link\]](#)
- [13] Jennifer Andreoli-Fang, John T. Chapman, “Mobile-aware scheduling for low latency backhaul over DOCSIS,” *Proc. of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, Montreal, Oct 2017. [\[link\]](#)
- [14] Jennifer Andreoli-Fang, John T. Chapman, “Latency reduction for mobile backhaul over DOCSIS through pipelining,” *Proc. of IEEE Globecom, Singapore*, Dec 2017. [\[link\]](#)

- [15] John T. Chapman, Jennifer Andreoli-Fang, Michel Chavin, Elias Chavarria Reyes, Zheng Lu, Dantong Liu, Joey Padden, Alon Bernstein, “Low latency techniques for mobile backhaul over DOCSIS,” *Proc. of IEEE Wireless Communication and Networking Conference (WCNC)*, Barcelona, April 2018. [[link](#)]
- [16] Jennifer Andreoli-Fang, John T. Chapman, “Synchronization for Mobile Backhaul over DOCSIS,” *SCTE Fall Technical Forum*, October, 2017. [[link](#)]
- [17] John T. Chapman, et. al., “The DOCSIS Timing Protocol (DTP), Generating precision timing services from a DOCSIS system,” *INTX/SCTE Spring Technical Forum*, 2011. [[link](#)]
- [18] Elias Chavarria Reyes, John T. Chapman, “How DOCSIS Time Protocol makes the SYNC Specification Tick,” *SCTE Fall Technical Forum*, Oct, 2020. [[link](#)]
- [19] “Synchronization Techniques for DOCSIS Technology Specification,” CM-SP-SYNC, CableLabs. [[link](#)]
- [20] John T. Chapman, “Remote PHY for Converged DOCSIS, Video and OOB”, *NCTA/SCTE Technical Forum*, Los Angeles, Jun, 2014. [[link](#)][[link2](#)]
- [21] John T. Chapman, “DOCSIS Remote PHY”, *SCTE Cable-Tec Expo Fall Technical Forum*, Atlanta, GA, Oct 2013.
- [22] “Remote PHY Specifications for DAA”, CM-SP-R-PHY, CableLabs. [[link](#)]
- [23] “nFAPI and FAPI specifications”, Small Cell Forum, May, 2017. [[link](#)]
- [24] Pawel Sowinski, Andy Smith & Tong Liu, “Remote PHY 2.0: The Next Steps For Remote PHY Technology”, *SCTE Cable-Tec Expo Fall Technical Forum*, New Orleans, Sep, 2019. [[link](#)]
- [25] Tong Liu, John T. Chapman, “Remote PHY with Remote Upstream Scheduler”, *SCTE Cable-Tec Expo Fall Technical Forum*, New Orleans, Sep, 2019. [[link](#)]
- [26] “Study on new radio access technology: Radio access architecture and interfaces,” 3GPP TR 38.801 Table A-1.
- [27] Dave Morley, John T. Chapman, Damian Poltz, Jennifer Andreoli-Fang, et. al., “Innovations in 5G Backhaul Technologies: IAB, HFC, and Fiber,” 5G Americas, June 2020. [[link](#)]

# **Collecting Smart City IoT Data to Generate Actionable Insights**

A Technical Paper prepared for SCTE•ISBE by

**Parmjit Dhillon**

Director Wireless R&D  
Charter Communications Inc  
6360 S. Fiddlers Green Circle, Greenwood Village, CO 80111  
(303) 793-4465  
Parmjit.Dhillon@charter.com

**Mohamed Daoud**

Principal Engineer  
Charter Communications Inc  
6360 S. Fiddlers Green Circle, Greenwood Village, CO 80111  
(720) 699-5077  
Mohamed.Daoud@charter.com



# Table of Contents

Title	Page Number
1. Abstract .....	3
2. Introduction.....	3
2.1. Data and the smart city .....	3
2.2. Why do we need smarter cities? .....	3
3. Proof of concept setup .....	4
3.1. Connectivity architecture .....	4
3.2. Hardware setup .....	4
3.3. Smart light controller .....	4
3.4. Camera.....	5
3.5. Pedestrian counter camera .....	5
3.6. Connectivity.....	5
3.7. WiFi .....	5
3.8. Weather station and air quality sensor.....	5
3.9. Emergency button .....	5
3.10. Digital signage and speaker.....	5
3.11. Smart intersection proof of concept .....	6
4. Data collection and analysis.....	7
4.1. Smart light energy usage .....	7
4.2. Camera.....	10
4.3. Smart intersection .....	12
4.4. Environmental .....	14
5. Conclusion.....	17
6. Acknowledgements .....	17
Abbreviations .....	17
Bibliography & References.....	17

## List of Figures

Title	Page Number
Figure 1 – Hardware Setup .....	4
Figure 10 – Loitering Alarm.....	11
Figure 12 – Computer Vision for Smart Intersection.....	12
Figure 13 – Pedestrian Count .....	13
Figure 14 – Monitoring Pedestrian Movement at Smart Intersection.....	14
Figure 15 - Weather .....	14
Figure 16 – Air Quality .....	15
Figure 17 – Air Quality (PM2.5) .....	16
Figure 18 – Insights from IoT Data .....	16

## List of Tables

Title	Page Number
Table 1 – Energy usage per day (Dimming vs No Dimming).....	9

## 1. Abstract

Data is the new oil driving the digital world. There is a massive amount of data being generated as more and more cities around the world roll out smart city applications and deploy IoT sensors. Managing, analyzing and correlating data from multiple sources can help cities in many ways, including lowering operational costs, creating environmental sustainability, increasing citizen engagement, enhancing healthcare, improving public safety and providing an overall improvement in the quality of life for its citizens.

Spectrum is engaged in a smart city proof of concept (POC) in St. Petersburg, FL, in collaboration with the University of South Florida St. Petersburg campus, US Ignite and the St. Petersburg Innovation District. Several IoT sensors and use cases involving pedestrian safety, video analytics, edge computing and smart street light management are currently being tested.

This paper discusses the uses cases and several types of sensor data collected in the St. Petersburg POC project. Details include the collection, storage, visualization and analysis of IoT sensor data.

## 2. Introduction

Data-driven decisions are made more than ever before in every domain today. Cities can also benefit from IoT sensor data to draw insights that can improve the quality of life of its citizens, allow cities to make better decisions on infrastructure planning and can help lower operating costs by efficiently using resources to deliver services to its citizens. In this paper, we share some examples of smart city use cases, technology and the methodology used to collect and analyze a broad array of IoT data from the proof of concept project implementation at the city of St. Petersburg, FL in collaboration with US Ignite, the St. Petersburg Innovation District, the University of South Florida St. Petersburg and Spectrum Enterprise.

### 2.1. Data and the smart city

Data is ultimately what makes a city smart. But it not just the collection of data through the deployment of sensor technology, it's also the ability to analyze and use the data in near real time to make informed decisions. When it comes to data, privacy is, of course, very important, especially if the data collected has personal information (PI). Such data should be handled as per the agreement with the city, and it must also meet regulatory compliance.

### 2.2. Why do we need smarter cities?

As cities around the world grow at an unprecedented rate, it is putting a strain on the environment and on cities' resources and services such as water, electricity, parking, waste management and many more. The gap between demand and supply of city resources is growing, affecting the quality of life of citizens. One way to close this gap is to use a city's limited resources more efficiently by deploying smart city applications. A forecast from the International Data Corporation (IDC) Worldwide Smart Cities Spending Guide shows global spending on smart cities initiatives will total nearly \$124 billion this year, an increase of 18.9% over 2019.[1]

### 3. Proof of concept setup

In the POC deployed in St. Petersburg, FL, we are using several IoT sensors, edge computing, computer vision, wired and wireless connectivity, cloud services for data storage, analysis and visualization to help the city improve pedestrian safety, lower energy costs for street lights and improve citizen engagement and safety.

#### 3.1. Connectivity architecture

In this POC, we have installed four smart light poles. Each light pole has dual connectivity. Fiber as a primary and 4G LTE as a secondary connectivity to backhaul the data to the cloud from all the sensors.

#### 3.2. Hardware setup

The fiber switch inside the light pole has multiple Power over Ethernet (PoE) ports providing power and data connectivity to a camera, WiFi, emergency button, digital banner, weather and air quality station as shown in Figure 1.

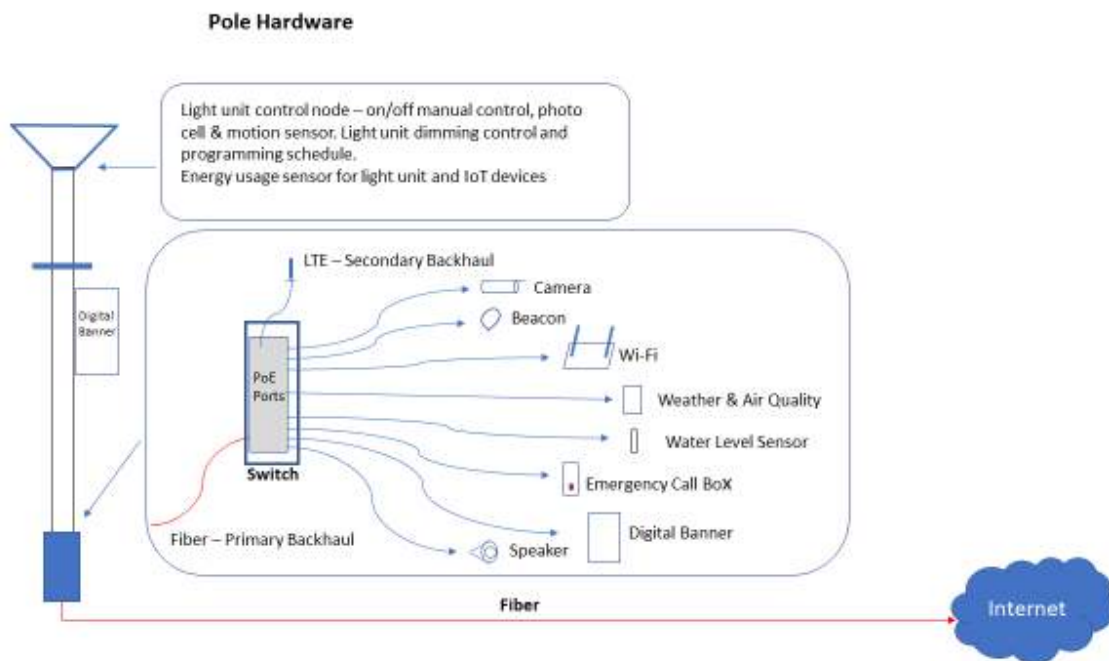


Figure 1 – Hardware Setup

#### 3.3. Smart light controller

The smart light pole uses energy-efficient LED lightbulbs. The smart light can be controlled via an easy-to-use user interface which provides a broad range of capabilities and functionality. The light can be remotely operated to activate at certain hours. Schedules for events can be set up or the light can automatically turn on when the photocell measures the ambient light level below a certain threshold.

Additionally, the light allows for integrated notification capabilities to provide updates or alerts to the operator when issues occur, delivering the city with a more efficient overall operating model.

### **3.4. Camera**

Each smart light pole is fitted with a camera that can be used for a variety of use cases including situational awareness and public safety. The camera has several alert features, e.g., guard zone, virtual fence, loitering alert and motion sensors to enhance public safety and protect city assets. We are using two types of cameras on the smart light pole. The Pan Tilt Zoom (PTZ) camera and fixed view camera with a 180-degree viewing angle. The power and data connectivity to the camera is provided by PoE switch as shown in Figure 1.

### **3.5. Pedestrian counter camera**

The pedestrian counter camera on the smart light pole is pointed downwards overlooking the sidewalk and is used to monitor the number of pedestrians and their direction of movement.

### **3.6. Connectivity**

Each smart light is also enabled with fiber connectivity to support a broad range of the features and functionality embedded in the light structure, from camera technology to integration of the smart intersection.

### **3.7. WiFi**

Each light is fitted with WiFi AP (802.11ac) with an external antenna mounted at the top of the light pole.

### **3.8. Weather station and air quality sensor**

The weather and air quality sensor is mounted on the top of the smart light pole and can collect various environmental parameters at specified intervals.

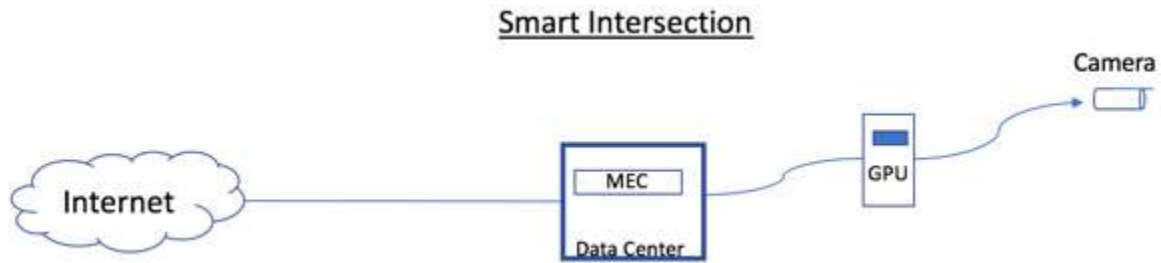
### **3.9. Emergency button**

Each smart light pole has an emergency voice over Internet protocol (VoIP) call box and routes the call to the iPBX over the fiber backhaul. The emergency call box is programmed to dial the local police emergency number. Each emergency call box hardware address and location is registered in the call manager, which enables the first responders to identify call origination location.

### **3.10. Digital signage and speaker**

The digital banner and the speaker can be remotely managed via a cloud-based dashboard. The display image and audio file can be sent over the network to be played at the smart light pole.

### 3.11. Smart intersection proof of concept



**Figure 2 – Smart Intersection System**

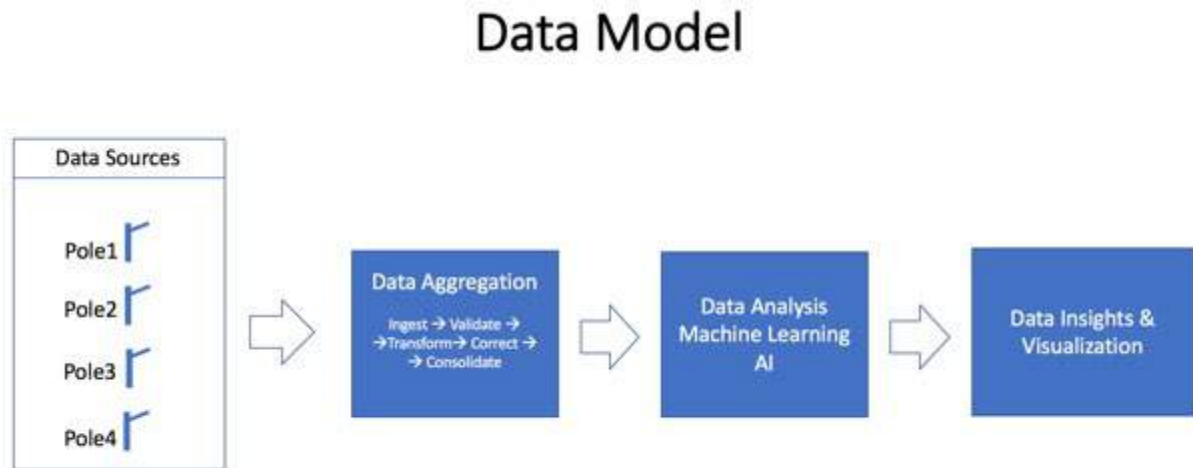
The smart intersection POC solution as shown in Figure 2 is designed to increase pedestrian safety at an intersection near the university. This POC requires ultra-low latency and uses edge computing technology with high computing power offered by the graphical processing unit (GPU) to process computer vision. There are three cameras mounted at the intersection as shown in Figure 3.



**Figure 3 – Smart Intersection Camera Zones**

## 4. Data collection and analysis

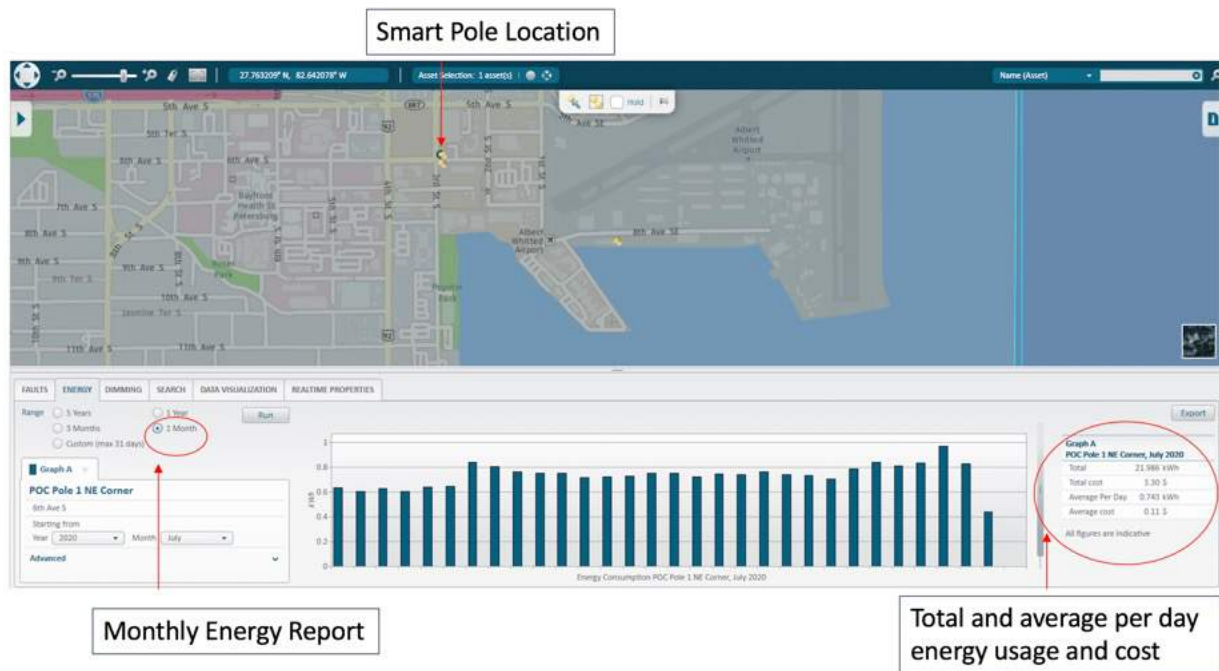
In this section, we will discuss the sensor data collection and flow as shown in Figure 4 below. Then we will highlight the data processing capabilities and driving insights in three examples: smart light energy usage, video analytics of surveillance camera and the smart intersection.



**Figure 4 – The Data Model**

### 4.1. Smart light energy usage

The smart light offers significant savings in energy expenses to a city, and we wanted to analyze ways to lower energy usage by using different settings on the unit. A LED lightbulb draws almost half of the energy as compared to High-Pressure Sodium (HPS) lightbulb. These savings can be further increased by using IoT sensors to control the light unit.

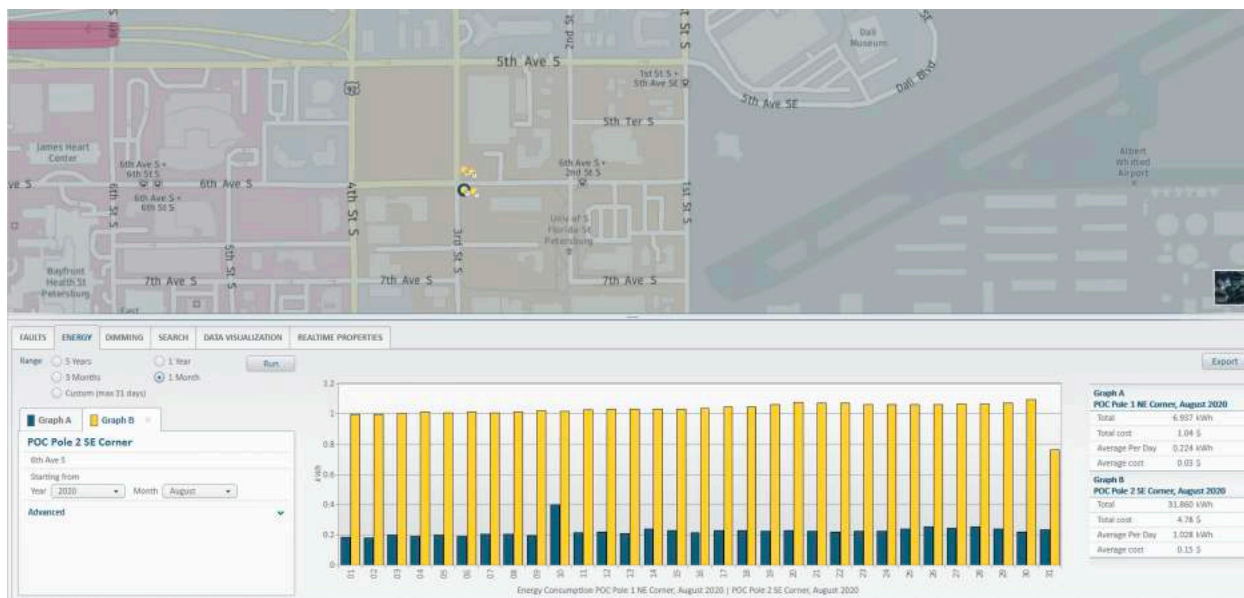


**Figure 5 - Dashboard**

The dashboard for the smart light as shown in figure 5 provides a means to view, control and manage the smart light remotely, thus increasing operational efficiency and lowering operating costs. The operating rules for the smart light, such as turn on and turn off schedule and managing the brightness level can be set remotely. We tested the two different settings on the smart light pole:

1. With dimming activated - in this setting the brightness changes when pedestrian movement is detected (via motion sensor/camera)
2. Without dimming - the light stays at full brightness level at all times.

We recorded the lowest energy usage with the dimming control activated as shown in Figure 6. The motion sensor increased the brightness of the light only when pedestrian activity was detected. When no activity was detected, the light pole stayed in dim state and provided maximum energy savings and a lower carbon footprint.




**Figure 6 – Energy Usage (Dimming vs No Dimming)**

**Table 1 – Energy usage per day (Dimming vs No Dimming).**

Smart Light Configuration	Average Energy Per Day (Watt Hour)
Without Dimming	1028
With Dimming Activated	224

### Energy usage during winter and summer months

We studied the variation in energy usage during the winter and summer months as shown in Figure 7. This information is useful for preparing for the load on the energy grid for each month.



**Figure 7 – Energy Usage in Winter and Summer Months**



Insights from alerts: The smart light system monitors many parameters and can send alerts when it detects an unusual reading from the sensor as shown in Figure 8. These insights can point to expected fluctuations on the power grid or high usage of the camera as result of pedestrian activity in the area, emergency button or WiFi.

**Warning:** The measured power consumption is significantly higher than specified nominal power.

Severity	ID	Category	Street	Name	Location Comment	First Reported On	Last Reported On	Created on	Detail	Component ID	Component Kind	Component Model	Is Open
Warning	46	Hardware failure	DEFAULT (World)	Brightline Pole 1 Devices		6/19/2020 1:43:59 PM	6/19/2020 1:43:59 PM	6/19/2020 4:29:46 PM	The measured power consumption 24.99W is significantly higher than specified nominal power 20.75W.	137	Communications Node	LLC7260	False
Error	50	Info	DEFAULT (World)	Pole 2 Devices		6/24/2020 10:02:14 AM	6/24/2020 10:02:14 AM	6/24/2020 10:02:20 AM	Grid unexpectedly switched on during day light.	136	Communications Node	LLC7260	False
Error	51	Info	7th Ave (World, St Pete South) (Energy)	Warehouses Devices		6/24/2020 12:22:42 PM	6/24/2020 12:22:42 PM	6/24/2020 12:22:47 PM	Grid unexpectedly switched on during day light.	12	Communications Node	LLC7260	False

**Figure 8 – Alarms and Alerts**

## 4.2. Camera

The camera is one of the most effective tools in the smart city tool kit for enhancing public safety and protecting city assets. As an example, many cities have some renovation and new development activities in progress. At these construction sites, expensive construction equipment and supplies are stored. After the work shift, when the workers leave the site, there is a possibility of theft or vandalism. Each year around \$300 million to \$1 billion worth of construction equipment is stolen and only less than 25% is recovered. Florida is ranked third in the nation for construction equipment theft.[2] We tested several camera alarms, e.g., virtual fence alarm, guard zone alarm and loitering alarm as shown in Figures 9, 10 and 11, at a construction site across the street from the proof of concept light pole.

A virtual fence has been useful to monitor unauthorized access at the construction site.



**Figure 9 – Virtual Fence**



Figure 20 – Loitering Alarm



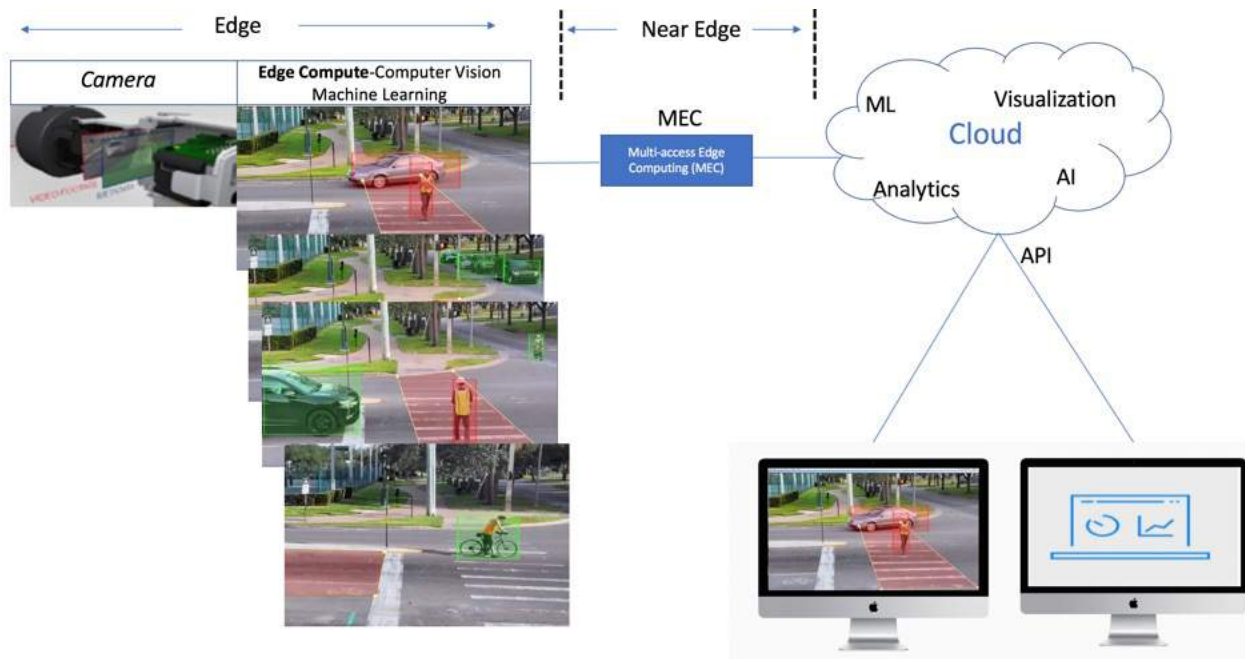
Figure 11 – Guard zone alert

### 4.3. Smart intersection

The goal of the Smart Intersection proof of concept is to improve pedestrian safety. The State of Florida has nine of the top 20 least safe cities in the U.S. for pedestrians and is also considered the most dangerous state for the pedestrians.[3] As per the National Highway Transport Safety Administration, pedestrian fatalities in crashes increased 53 percent in the last decade (2009 to 2018), with the pedestrians' share of traffic fatalities increasing 42 percent, from 12 to 17 percent.[4]

One initiative cities like St. Petersburg are working on is to bring the number of fatalities or collisions in an intersection to zero. This effort is called Vision Zero. Data collected from road sensors and cameras can help cities better understand everything from crosswalk signal timing to traffic patterns that may lead to these collisions.

The Smart Intersection proof of concept uses computer vision and edge computing to detect the presence of a pedestrian as shown in Figures 12 and 14. The camera's technology can be used to set the zone within the view of the camera that needs to be monitored for pedestrian activity. The anonymous pedestrian and vehicular traffic data is stored in the cloud for further analysis, planning and design of components of the intersection, including stop signs, traffic/pedestrian light timing, crosswalks and sidewalks to improve pedestrian safety.



**Figure 32 – Computer Vision for Smart Intersection**


We are currently collecting the following metadata as shown in Figure 13 to assist St. Petersburg with the design of the intersection that would improve pedestrian safety:



- Number of pedestrians crossing the crosswalk
- Direction of pedestrian movement e.g., north, south, east, west

	Last Hour	Daily	Monthly
North to South	11	242	7260
South to North	22	484	14520
East to West	89	1958	58740
West to East	113	2486	74580
Totals	235	5170	155100

**Figure 43 – Pedestrian Count**



**Figure 54 – Monitoring Pedestrian Movement at Smart Intersection**

#### **4.4. Environmental**

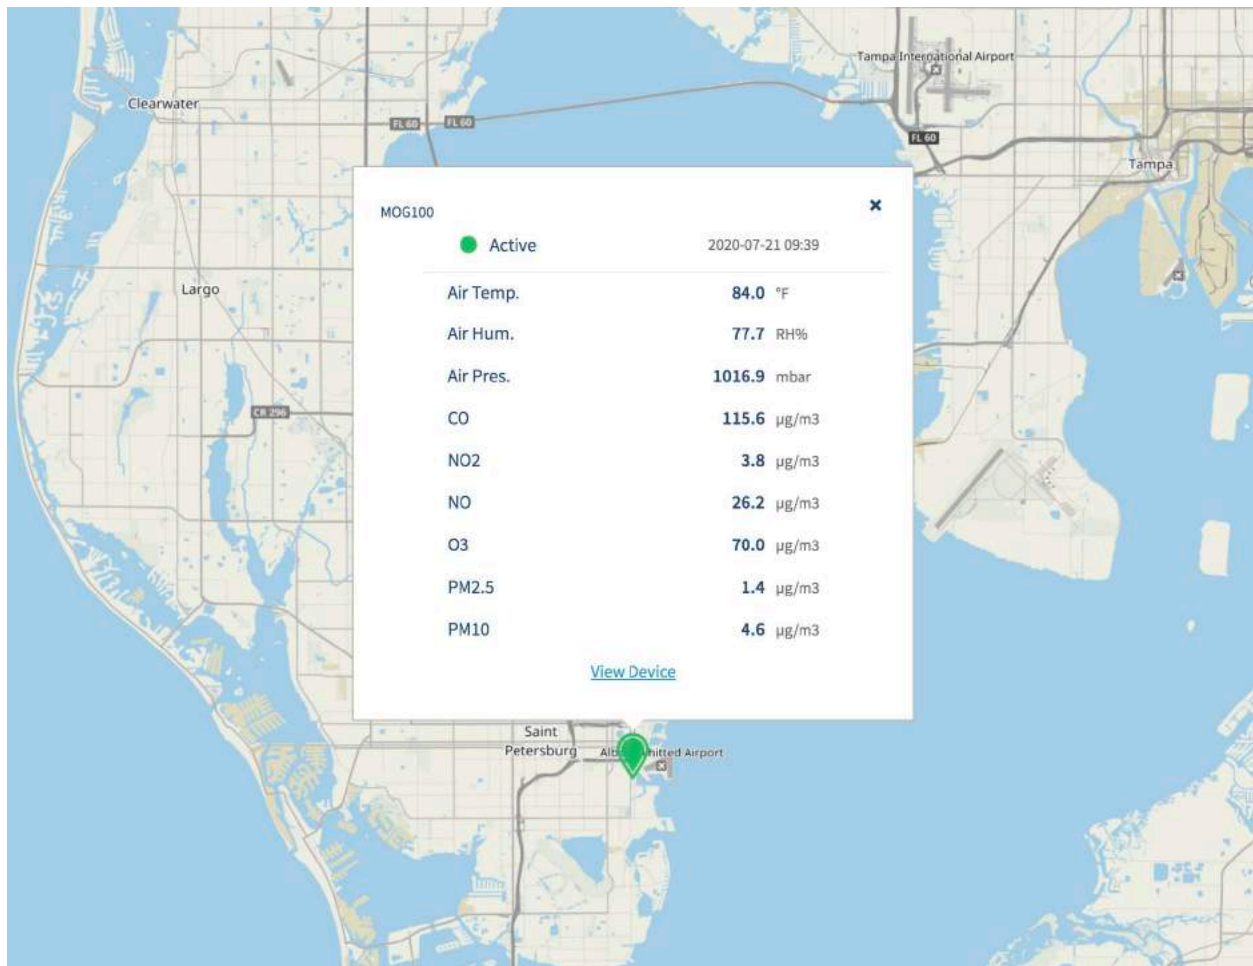
Florida is home to some extreme weather conditions including hurricanes, extremely hot temperatures, humidity, fog and a high number of lightning strikes from thunderstorms. Florida leads the country in both the number of fatal lightning hits and the density of strikes, with just over 20 flashes per square mile every year.[5]



**Figure 65 - Weather**

Weather can have a big impact on citizen safety, city operations and its economy. Cities and businesses today not only need a general weather forecast for the whole city, but a micro-forecast for each neighborhood. The weather sensors are one of the largest sources of IoT data and help in micro forecasting of weather. Data from local weather sensors can micro-forecast weather that can help utility companies predict the energy load for each part of the city. Cities can predict and monitor flooding in each neighborhood as shown in Figure 15 and prepare the water pumps; detect a lightning strike at a particular location and alert the public; and detect fog on a certain part of the highway and warn motorists.[6]

Air pollution is responsible for five million premature deaths each year and nine out of ten people around the world breathe unhealthy air.[7] The sensors installed on the light pole in the POC collect several weather and air quality parameters as shown in Figures 16 and 17 and provide the city with interesting and pertinent insights on air quality at the location of the sensor.



**Figure 76 – Air Quality**

Following are some of the parameters collected by this sensor:

**Table 2 – Weather and Air Quality**

Weather	Air Quality
<ul style="list-style-type: none"> <li>• Temperature</li> <li>• Humidity</li> <li>• Dew Point</li> <li>• Air Pressure</li> <li>• Wind Direction and Speed</li> <li>• Precipitation Intensity</li> <li>• Radiation</li> <li>• Lightning</li> </ul>	<ul style="list-style-type: none"> <li>• Carbon Monoxide</li> <li>• Nitrogen Dioxide</li> <li>• Nitric Oxide</li> <li>• Ozone</li> <li>• PM 1</li> <li>• PM 10</li> <li>• PM 2.5</li> </ul>




Figure 87 – Air Quality (PM2.5)

The data from all the sensors is collected and sent to the cloud where it is analyzed to draw actionable insights, as shown in Figure 18

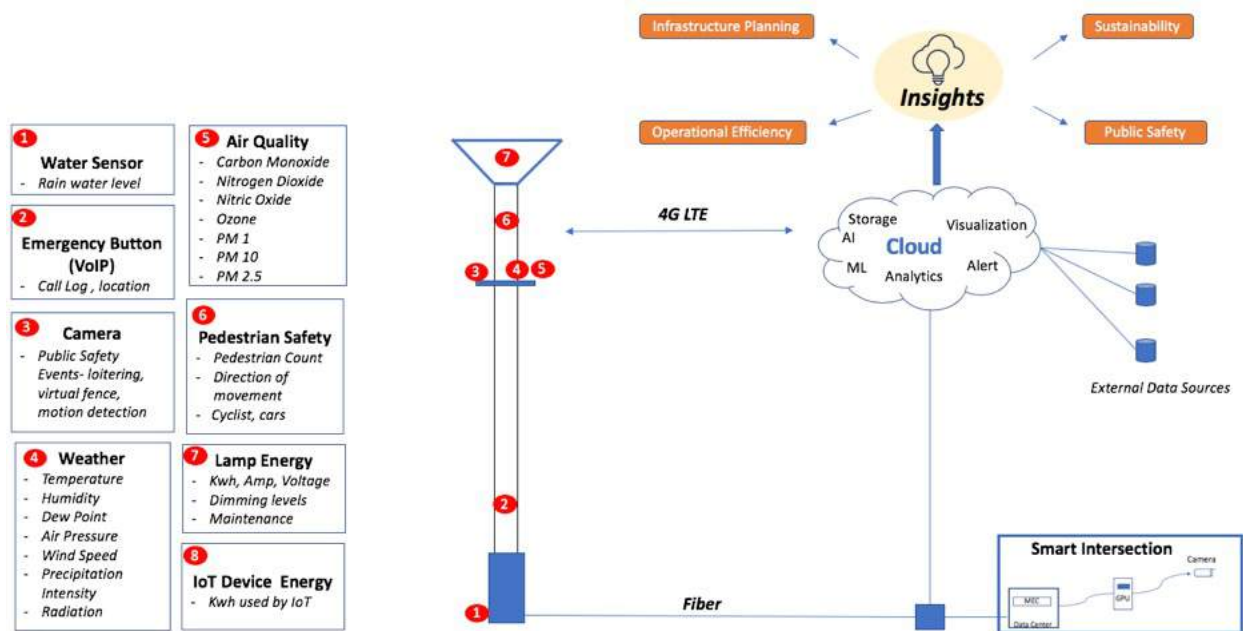


Figure 98 – Insights from IoT Data

## 5. Conclusion

In this proof of concept, we have demonstrated how much data can be collected and potentially analyzed from a few key IoT sensors.

As we collect more data and datasets get richer over time, we expect to draw more insights. Using the richer data sets, we can train the computer vision model and capture additional metadata to make our roads safer, predict and manage energy usage to lower costs and reduce the carbon footprint, and improve public safety and citizen engagement.

Installation of IoT hardware requires permits and inspections from many city departments before the sensors can be activated. When selecting a location for a POC, pick multiple locations that represent different sections of the city and also future-proof the POC to add scale and capability.

Close collaboration with academia, local communities and the government is important for the successful deployment of smart city solutions. Each city is unique; therefore, smart city solutions and use cases should be customized. The cable companies, with its vast network infrastructure and strong public-private partnership, are well positioned to lead the way in offering smart city solutions.

## 6. Acknowledgements

The authors would like to acknowledge Alison Barlow from the Innovation District at the city of St. Petersburg, FL and representatives of the University of South Florida St. Petersburg campus, US Ignite and St. Petersburg, FL, for the support and guidance during the execution of the POC.

## Abbreviations

IP-PBX	Internet Protocol Private Branch Exchange
MEC	Multi-access Edge Computing
POC	Proof of Concept
PoE	Power over Ethernet
IoT	Internet of Things

## Bibliography & References

*Environmental Defense Fund*. (n.d.). Retrieved 07 21, 2020, from <https://www.edf.org/health/health-impacts-air-pollution#:~:text=Air%20pollution%20is%20now%20the,AIDS%2C%20tuberculosis%20and%20malaria%20combined.>

*IDC*. (2019, June 25). (IDC) Retrieved July 5, 2020, from <https://www.idc.com/getdoc.jsp?containerId=prUS45303119>



- Johnson, C. (2019, January 29). *Report: Florida still the most dangerous state for pedestrians*. Retrieved July 05, 2020, from Tampa Bay Times : <https://www.tampabay.com/transportation/report-florida-still-the-most-dangerous-state-for-pedestrians-20190123/>
- Jones, K. (2017, 07 30). *Construct connect*. Retrieved 07 21, 2020, from <https://www.constructconnect.com/blog/high-cost-construction-equipment-theft>
- National Highway Transport Safety Administration. (2019, September). Retrieved July 05, 2020, from <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812822>
- Tampa Bay Times. (2019, Feb 7). Retrieved July 05, 2020, from [https://www.tampabay.com/weather/hey-florida-lightning-strike-capital-of-the-us-say-goodbye-to-dr-lightning-20190207/#:~:text=goodbye%20to%20Dr.,Lightning,per%20square%20mile%20every%20year.&text=Florida%20is%20the%20lightning%20strike,more%20intimately%](https://www.tampabay.com/weather/hey-florida-lightning-strike-capital-of-the-us-say-goodbye-to-dr-lightning-20190207/#:~:text=goodbye%20to%20Dr.,Lightning,per%20square%20mile%20every%20year.&text=Florida%20is%20the%20lightning%20strike,more%20intimately%20)
- Williams, D. P. (n.d.). Retrieved July 05, 2020, from <https://meetingoftheminds.org/smart-cities-weather-22100>

# Optimizing the 10G Transition to Full-Duplex DOCSIS® 4.0

A Technical Paper prepared for SCTE•ISBE by

**Richard S Prodan, Ph.D.**  
Engineering Fellow  
Comcast Cable  
1401 Wynkoop Street #300  
720-512-3742  
rich\_prodan@comcast.com

# Table of Contents

Title	Page Number
1. Introduction.....	4
2. Cable Network Design .....	4
3. The 10G Transition Path to Full-Duplex DOCSIS.....	6
4. Preparing the Transition from Sub-Split to Mid-Split.....	6
5. Completing the Transition to Mid-Split or Upgrading to High-Split .....	13
6. Accomodating Legacy MoCA in 1.2 GHz System Upgrades .....	16
7. Upgrading to Full Duplex DOCSIS for the Node + 0 Architecture .....	22
8. Optimizing Full Duplex Architecture and Operation .....	26
8.1. FDX Modem Downstream Legacy Band Receive Power .....	26
8.2. FDX Modem FDX Band Receive Power .....	28
8.3. Signal-to-Echo Ratio Optimization .....	30
9. Conclusion.....	33
Abbreviations .....	35
Bibliography & References.....	36

## List of Figures

Title	Page Number
Figure 1 – Conventional Node + N network architecture .....	4
Figure 2 – RF Spectrum: Sub-Split STB Compatible – 750 MHz Plant.....	5
Figure 3 – RF Spectrum: Mid-Split – 750 MHz Plant.....	5
Figure 4 – RF Spectrum: High-Split – 1 GHz Plant .....	5
Figure 5 – Remote PHY Architecture.....	6
Figure 6 – In-home video and data connection to the cable plant.....	7
Figure 7 – Spurious emissions below diplex filter cut-off.....	7
Figure 8 – Percentage of STBs $\leq$ CIR and CNIR.....	10
Figure 9 – Node Sub-Split Video Transmit Power Boost in the Mid-Split Band .....	12
Figure 10 – High-Split CM ACI into Neighbor Legacy STB .....	14
Figure 11 – Percentage of STBs $\leq$ CACIR .....	16
Figure 12 – MoCA Non-Overlapping Channel & DOCSIS 3.1 Downstream Frequency Limit.....	17
Figure 13 – MoCA 2.0 Field Trial Results of Percentage of Outlets vs. MAC Rate.....	18
Figure 14 – Analysis of MoCA Susceptibility into DOCSIS 3.1.....	20
Figure 15 – MoCA 2.0 Transmitter to DOCSIS 3.1 Receiver RF Isolation .....	21
Figure 16 – MoCA 2.0 in 1.1 mode Transmitter to DOCSIS 3.1 Receiver RF Isolation .....	21
Figure 17 – FDX Interference Analysis Model in the Node + 0 Network .....	22
Figure 18 – Configurable FDX Allocated Spectrum Bandwidths .....	23
Figure 19 – RF Spectrum: FDX with Sub-Split, Mid-Split and High-Split – 1.1 GHz Plant.....	24
Figure 20 – RF Spectrum: FDX with Sub-Split, Mid-Split and High-Split – 1.2 GHz Plant.....	25
Figure 21 – RF Spectrum: Full FDX with Sub-Split, Mid-Split and High-Split – 1.2 GHz Plant (Compliant with DOCSIS 4.0 Spec).....	25
Figure 22 – RF Spectrum: Full FDX with Sub-Split, Mid-Split and High-Split – 1.2 GHz Plant (Non-Compliant with DOCSIS 4.0 Spec) .....	26

Figure 23 – Specification Topic, CPE Architecture: Option A vs Option B Discussion.....	27
Figure 24 – CM Architecture with Filter Response Avoiding Legacy Downstream Loss.....	27
Figure 25 – FDX Remote PHY Node Functional Block Diagram.....	28
Figure 26 – FDX Node RF Level with FDX Band Downstream Level Boosted +7 dB.....	29
Figure 27 – FDX Node RF Level with FDX Band Downstream Level Boosted +4 dB.....	29
Figure 28 – Upstream Node Signal/Interference Ratio vs FDX Downstream Band PSD Boost (Top: +4 dB Boost, Middle: No Boost – Linear, Bottom: -4 dB Attenuation) .....	31
Figure 29 – Downstream CM Signal/Interference Ratio vs FDX Downstream Band PSD Boost (Top: +4 dB Boost, Middle: No Boost – Linear, Bottom: -4 dB Attenuation) .....	32

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Mid-Split Spurious Emissions of Sub-Split Upstream Carriers .....	8
Table 2 – Set-Top Box Video Interference Threshold .....	8
Table 3 – Analysis of Percentage of STBs Below Threshold SNR.....	11
Table 4 – Analysis of Percentage of STBs Below Threshold SNR with 6 dB Power Boost .....	12
Table 5 – High-Split CM Carrier/ACI Ratio at STB Impaired Video Threshold.....	14
Table 6 – Analysis of Percentage of STBs Below Video Threshold CACIR.....	15
Table 7 – MoCA Extended Band D Frequency Plan .....	18
Table 8 – DOCSIS 3.1 RF Specifications .....	19
Table 9 – MoCA RF Specifications .....	19

# 1. Introduction

Comcast has been deploying optical fiber deeper into its infrastructure, in “Node + 0” /, “N+0” mid-split configurations to support the ever-increasing demand for higher speeds and throughput to customers. The first step was the transition from traditional Node + N sub-split networks, with 5 to 42 MHz upstream and 1002 MHz downstream bandwidth, to current mid-split Node + 0 networks with 5 to 85 MHz upstream and 1218 MHz downstream bandwidth.

Future Node + 0 networks will transition to Full-Duplex (FDX) DOCSIS, to significantly increase the upstream bandwidth to multi-gigabit speeds with high throughput. This will not be accomplished in one giant leap, but rather a transition from mid-split to high-split networks. The first step will be a shift to a 5 to 204 MHz upstream and 1218 MHz downstream bandwidth, supporting a 1 Gbps upstream tier and multi-Gbps downstream. As demand for multi-gigabit services grows, FDX can be deployed, with overlapping bidirectional spectrum from 108 to 204 MHz that eventually increases up to the full 108 to 684 MHz FDX limit plus legacy DOCSIS 3.0 to the 1002 MHz limit and legacy DOCSIS 3.1 to the 1218 MHz limit.

This paper will discuss considerations that need to be addressed in the transition. Scenarios will be described, as well as some node and modem architectures that provide flexible signal support to optimize the end-to-end performance, toward 10G Full-Duplex operation of the network.

## 2. Cable Network Design

Traditional hybrid fiber/coax (HFC) cable distribution networks were built as-tree-and-branch networks consisting of a fiber node connecting multiple cascaded amplifier coax cable sections. Each section connects to a series of multiport taps, transmitting signal to and receiving signals from drop cables to customer premise equipment (CPE.) An example of one coax branch of a conventional Node + N HFC architecture is shown in Figure 1. The node span contains multiple amplifier spans, each with multiple taps between amplifiers.






Figure 1 – Conventional Node + N network architecture


This conventional architecture provides two-way signal transmission on separate spectral bands using frequency division duplex (FDD) operation. Each amplifier section uses diplex filtering to separate upstream transmissions, toward the node, in the narrower lower frequency band (typically 42 MHz for sub-split, 85 MHz for mid-split, and 204 MHz for high-split systems) from downstream transmissions, from the node, in the much wider upper frequency band (up to 1.2 GHz). RF spectrum examples for sub-split, mid-split, and high-split systems are shown in Figure 2, Figure 3, and Figure 4 respectively.



**Figure 2 – RF Spectrum: Sub-Split STB Compatible – 750 MHz Plant**



**Figure 3 – RF Spectrum: Mid-Split – 750 MHz Plant**




**Figure 4 – RF Spectrum: High-Split – 1 GHz Plant**

Such diplex filtering prevents two-way transmission within the same spectral band. Each multiport tap contains a directional coupler that diverts a portion of the downstream signal to the drops connected to the tap ports, and injects the upstream signals present on the tap ports toward the node. The directivity of the directional coupler prevents upstream signals from propagating in the downstream direction or from diverting to other drops upstream from that tap port.

Comcast has been deepening its deployment of fiber via Node + 0 mid-split systems, which separate the CMTS core functions that connect to Remote PHY (R-PHY) nodes using a Distributed Access Architecture (DAA). The Remote PHY Architecture shown in Figure 5 moves the RF modulation and demodulation closer to the customer, either in a node in the HFC plant or in a hub location.

The Media Access Control (MAC) layer processing is in the headend, where a CCAP-Core is responsible for all MAC-layer processing and multiplexing. The Remote PHY Device (RPD) performs the PHY-layer functions and transmits the PHY-processed RF signals on subscriber ports connected to coax segments closer to the customer. The RPD is connected to the CCAP-core via a digital Ethernet connection over digital fiber optics.



**Figure 5 – Remote PHY Architecture**

The digital replacement of conventional analog fiber optics and elimination of cascaded amplifiers in the tapped coax cable spans from the node, as well as the reduced distance to the customer from the node, results in a substantial increase in delivered Signal to Noise Ratio (SNR.)

### 3. The 10G Transition Path to Full-Duplex DOCSIS

The deployment of FDX will follow an evolutionary path involving upgrading the fiber depth into the networks, in mid-split and/or high-split FDD operation, to limited FDX bandwidth operation. Then the amount of FDX spectrum will gradually expand as future demand for higher speed and throughput increases.

However these transition scenarios require management of both the legacy and new consumer premises equipment (CPE) and R-PHY node upgrades to support multiple generations of video and data services. Currently, Node + 0 mid-split systems supply video to legacy sub-split set-top boxes (STBs) in portions of the mid-split spectrum while simultaneously deploying new mid-split cable modems in the same customer premises.

As the transition progresses, the R-PHY node will need to be upgraded. New CMs will need to be added to support higher speed tiers with increased spectral efficiency. A description of the intermediate steps to support this 10G evolution, some problems that arise in supporting legacy CPE in upgraded systems, and solutions to mitigate these problems are presented.

### 4. Preparing the Transition from Sub-Split to Mid-Split

The 10G transition provides higher speeds and throughput with reduced latency. Getting there requires expanding capacity -- primarily the upstream, but also the downstream spectrum as demand increases. Currently, our cable networks are a mix of sub-split, business as usual (BAU) Node + N networks, and mid-split Node + 0 networks. Upgrading BAU topologies or taking fiber deeper to support mid-split spectrum can happen, in terms of DOCSIS 3.1-based cable modems (CMs) and gateways, by launching initially in sub-split mode. This supports legacy (sub-split) STBs while the network is being prepared for the transition to future mid-split operation. The mid-split CMs will be “seeded” into the existing sub-split networks, operating within the existing sub-split RF spectrum. All DOCSIS CMs receive downstream channels in the mid-split downstream frequency band above 108 MHz regardless of the upstream frequency split (42 or 85 MHz), except DOCSIS 3.1 high-split CMs (with upstream and downstream band edges at 204 and 258 MHz respectively).

All upgrade considerations for mid- and high-splits would do well to revisit the notoriously noisy and hostile spectral environment that is the upstream path. In particular, a mid-split CM transmitting upstream in the sub-split 5 to 42 MHz spectrum can produce spurious emissions in the 42 to 85 MHz band, below the mid-split cutoff of the CM's duplex filter. Also within that region, between 54-85 MHz, are STB downstream video channels prior to enabling mid-split operation. A customer equipped with a DOCSIS mid-split CM and a sub-split STB will be connected to the drop via an in-home splitter, as shown in Figure 6.

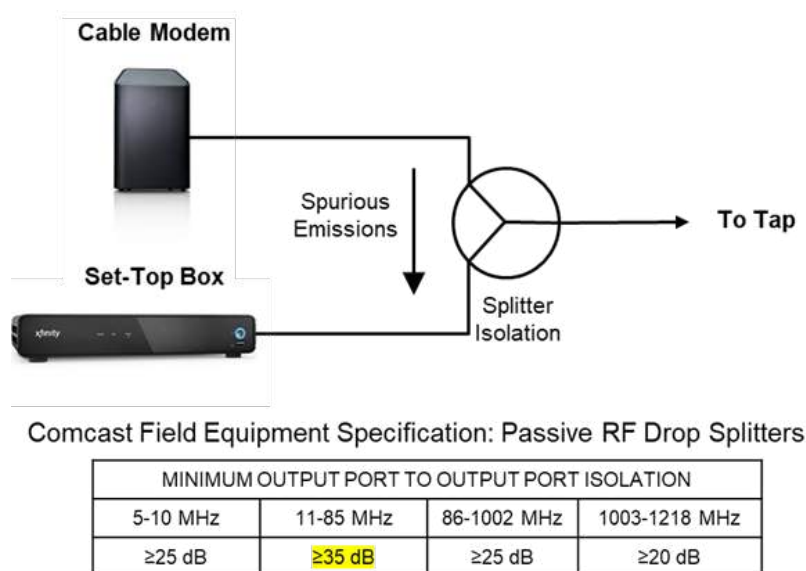


Figure 6 – In-home video and data connection to the cable plant

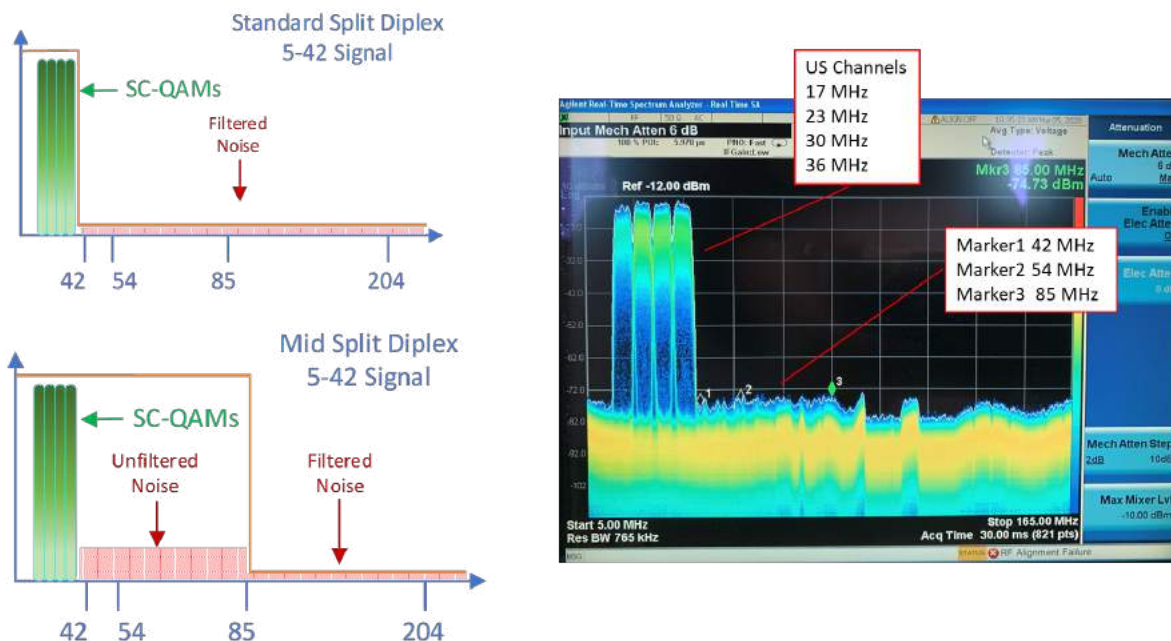


Figure 7 – Spurious emissions below duplex filter cut-off



Figure 7 shows spurious noise in the mid-split CM 5-42 MHz upstream transmission path that can interfere with an STB's downstream video reception between 54 to 85 MHz. Measurements of the spurious emissions levels and the reduced levels below the spec limit of -44 dBr are shown in Table 1. Notice that the highlighted maximum spurious emissions level reduction at the maximum transmit power per channel is at least 10 dB below the -44 dBr spec limit.

**Table 1 – Mid-Split Spurious Emissions of Sub-Split Upstream Carriers**

Device Model	US TX power per channel (dBmV/6.4 MHz)	Spurious power @54 MHz (dBmV/200 kHz)	Spurious Emissions (dBr)	Spurious Emissions Reduction from -44 dBr spec (dB)
XB7	53.56	-17.91	-56.42	12.42
XB7	47.31	-19.81	-52.07	8.07
XB7	43.47	-23.92	-52.34	8.34
XB6 v1	53.76	-16.55	-55.26	11.26
XB6 v1	47.43	-20.85	-53.23	9.23
XB6 v1	43.82	-25.27	-54.04	10.04
XB6 v2	52.56	-17.28	-54.79	10.79
XB6 v2	46.75	-20.81	-52.51	8.51
XB6 v2	43.07	-21.3	-49.32	5.32

**Table 2 – Set-Top Box Video Interference Threshold**

STB Make	STB Model	Frequency	Measured MER (dB)	STB reported SNR (dB)
ARRIS	XG1v3/AX013ANM	141 MHz	27.1	27.2
ARRIS	XG1v3/AX013ANM	147 MHz	27.5	27.0
ARRIS	XG1v3/AX013ANM	159 MHz	27.5	27.1
ARRIS	XG1v3/AX013ANM	165 MHz	27.1	27.1
ARRIS	XG1v3/AX013ANM	177 MHz	27.2	27.1
ARRIS	XG1v1/MX011ANM	141 MHz	27.2	27.0
ARRIS	XG1v1/MX011ANM	147 MHz	27.3	27.1
ARRIS	XG1v1/MX011ANM	159 MHz	27.4	27.0
ARRIS	XG1v1/MX011ANM	165 MHz	27.0	27.0
ARRIS	XG1v1/MX011ANM	177 MHz	27.1	27.2
Pace	XG2v2/PX022ANM	141 MHz	27.0	27.1
Pace	XG2v2/PX022ANM	147 MHz	27.5	27.2
Pace	XG2v2/PX022ANM	159 MHz	27.4	27.2
Pace	XG2v2/PX022ANM	165 MHz	27.0	26.9
Pace	XG2v2/PX022ANM	177 MHz	27.0	27.0
Pace	XG1v3/PX013ANM	141 MHz	27.3	27.0
Pace	XG1v3/PX013ANM	147 MHz	27.0	26.9

Pace	XG1v3/PX013ANM	159 MHz	27.0	27.1
Pace	XG1v3/PX013ANM	165 MHz	27.1	27.0
Pace	XG1v3/PX013ANM	177 MHz	27.2	27.1

The port-to-port isolation of the in-home splitter in Figure 6 will determine the attenuation of the coupled interference. The magnitude of the interference depends on each CM transmit level, and the STB receive level separated by the splitter isolation. If the resulting interference from a CM lowers the SNR of the STB's downstream receiver below the video channel SC-QAM threshold, then video artifacts, or a complete loss of reception of some channels below 85 MHz, may occur for that home. Measurements of the STB SC-QAM video threshold SNR levels below the SCTE 40 spec limit of 30 dB that exhibited "tiling" (macroblock freezing/errors in portions of the decompressed video) is shown in Table 2, for several deployed STB models at several frequencies. The video SC-QAM level was -5 dBmV/6 MHz for these measurements. The maximum SNR threshold level where video artifacts are observed is 27.5 dB, which is more than 2 dB below the SCTE 40 spec limit.

Field data was obtained for the following analysis using the above determined parameters of mid-split CM noise and spurious emissions interference into STBs in the same home. Over 21 thousand CMs in Comcast employees homes were polled using Comcast field analysis tools tabulating over 23 million CM MER, transmit and receive levels at multiple frequencies with multiple samples per frequency for each CM from more than 800 CMTS network locations. The following analysis shows the number of cases that could cause observable STB video artifacts are few.

The SNR metrics for spurious emissions interfering with the video SC-QAM signal across a two-way splitter are the carrier-to-interference ratio (CIR) and the carrier-to-noise-and-interference ratio (CNIR). The interference level, I, is the spurious emissions level calculated from the DOCSIS 3.1 PHY spec using the measured upstream transmit power at the CM F-connector. The background noise level is determined by the reported downstream CM MER. All such measurements are averaged for each CM MAC address.

The maximum allowed spurious emissions level of each CM is calculated per the DOCSIS 3.1 PHY spec. For a legacy Transmit Channel Set (TCS) with four 6.4 MHz SC-QAM channels (5.12 Msym/s), the 100% Grant spectrum is 25.6 MHz. The measurement bandwidth (Measurement BW or Interval) is 1.6 MHz.

The SpurFloor is defined as

$$\text{SpurFloor} = \max[-57 + 10 \cdot \log(100\% \text{ Grant Spectrum}/192 \text{ MHz}), -60] \text{ dBc}$$

Thus the SpurFloor for the 100% Grant bandwidth in the Measurement BW is -60 dBc, and the emissions level within the Measurement BW is:

$$\text{Spec in Interval (dBc)} = -60 + 10 \log[\text{Measurement BW}/(100\% \text{ Grant}/40)] = -56 \text{ dBc}$$

$$\text{Emissions Limit (dBr)} = \text{Spec in Interval} + 10 \log[100\% \text{ Grant}/\text{Measurement BW}] = -44 \text{ dBr}$$


Note that the measured emissions limit in Table 1 shows a 10 dB reduction in practice. The noise and spurious emissions power per 6 MHz channel for the total composite power TCP dBmV of all TCS channels is:

$$\text{TCP} - 56 \text{ dBc} + 10 \log[6/1.6] \text{ dBmV}/6 \text{ MHz}$$

The data analysis method is as follows:

Calculate the average MER, transmit and receive level for each modem over all upstream and downstream frequencies per MAC address.

Calculate the average emissions level for each average transmit level for each CM MAC address.



**Figure 8 – Percentage of STBs ≤ CIR and CNIR**

Calculate the STB downstream CIR for each average receive level and each CM computed average emissions level as:

$$\text{CIR} = \text{DS Average Receive} - \text{US Average Emissions} + \text{Splitter Isolation}$$

$$\text{CNIR} = -10 \log(10^{-\text{CIR}/10} + 10^{-\text{MER}/10}).$$

Sort the calculated STB CIR and CNIR levels in ascending order (plotted in Figure 8).

Calculate the percentage of CMs  $\leq$  each STB CIR and CNIR level as:

$$\% \text{CMs} \leq \text{CIR value} = \text{Number of CIR entries} \leq \text{CIR} / \text{Total Number of CIR values} * 100$$

$$\% \text{CMs} \leq \text{CNIR value} = \text{Number of CNIR entries} \leq \text{CNIR} / \text{Total Number of CNIR values} * 100$$


The parameters used and analysis results shown in Figure 8 are given in the following table:

**Table 3 – Analysis of Percentage of STBs Below Threshold SNR**


Specifications		Threshold Video CIR	
CPE Parameters		Total Downstream Band	
Splitter Isolation (dB):	35	% Set-Top Boxes with < 28 dB CIR	7.24%
Set-Top Box Threshold CNR (dB):	28	Set-Top Box CIR(dB):	28.0
Spurious Emissions Reduction (dB):	10	Threshold Video CNIR	
Downstream Boost 54 to 85 MHz (dB):	0	Total Downstream Band	
D3.1 Spurious Emissions Spec		% Set-Top Boxes with < 28 dB CNIR	7.71%
100% Grant (MHz):	25.6	Set-Top Box CNIR(dB):	28.0
Spur Floor (dBc):	-60		
Under-grant Hold Number of Users:	40		
Under-grant Hold BW (MHz):	0.64		
Measurement BW (MHz):	1.6		
Spec in Interval (dBc):	-56		
Emissions Limit (dBr):	-44		

The cases which have sufficient interference magnitude will require mitigation. The percentage of STBs below the SNR threshold for the specifications in Table 3, including background noise plus spurious emissions, is below 8%. This percentage can be further reduced which decreases the number requiring 85 MHz low pass filters on the CMs on the same splitter as the low SNR STBs. Mitigation strategies that limit the number of cases where interference-blocking low pass filters needed at the CM output are discussed next.

The lower frequency sub-split video channel RF levels from 54 to 85 MHz can be boosted up to 10 dB at the node with a de minimus (0.04 dB) increase in total composite power (TCP) as shown in Figure 9.



Total Composite Power(dBmV): 73.8



Total Composite Power(dBmV): 73.84

**Figure 9 – Node Sub-Split Video Transmit Power Boost in the Mid-Split Band**

The result of applying a 6 dB downstream power boost below 85 MHz is given in the following table:

**Table 4 – Analysis of Percentage of STBs Below Threshold SNR with 6 dB Power Boost**

Specifications	
CPE Parameters	
Splitter Isolation (dB):	35
Set-Top Box Threshold CNR (dB):	28
Spurious Emissions Reduction (dB):	10
Downstream Boost 54 to 85 MHz (dB):	6

D3.1 Spurious Emissions Spec	
100% Grant (MHz):	25.6
Spur Floor (dBc):	-60
Under-grant Hold Number of Users:	40
Under-grant Hold BW (MHz):	0.64
Measurement BW (MHz):	1.6
Spec in Interval (dBc):	-56
Emissions Limit (dBr):	-44

Threshold Video CIR	
Total Downstream Band	
% Set-Top Boxes with < 28 dB CIR	1.68%
Set-Top Box CIR(dB):	28.0

Threshold Video CNIR	
Total Downstream Band	
% Set-Top Boxes with < 28 dB CNIR	1.90%
Set-Top Box CNIR(dB):	28.0

The percentage of STBs below the SNR threshold for the specifications with 6 dB boosted downstream transmit power below 85 MHz in Table 4 including both background noise plus spurious emissions is now below 1.9%. This virtually eliminates the mid-split CM emissions into a sub-split STB interference problem without an increase in the transmitted total composite power of the node.

The initial deployment of mid-split CMs described above prior to enabling mid-split operation were limited to transmit in the sub-split band. Video signals remained in the sub-split downstream band with the lower band edge at 54 MHz. The downstream signals can be moved to occupy anywhere in the 108 to 1218 MHz band to transition to mid-split operation. These mid-split upgrades will move the video channels anywhere between 108 to 1002 MHz. This will prevent the spurious emissions interference problem previously discussed, as the video channels will be moved above the 85 MHz upstream cutoff frequency of the mid-split CM's duplex filter.

However, the sub-split STBs will encounter a different problem after the transition enabling mid-split operation or when upgrading to high-split operation. This is discussed in the next section.

## **5. Completing the Transition to Mid-Split or Upgrading to High-Split**

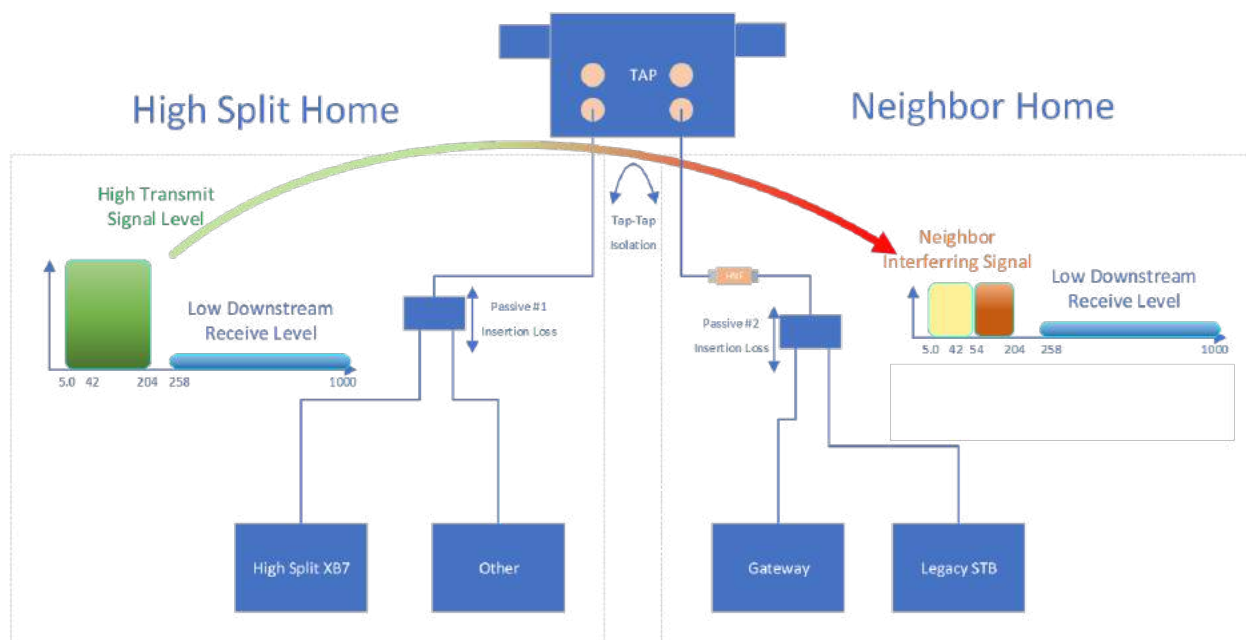
Node + 0 networks have been and continue to be deployed in a mid-split frequency division duplex (FDD) configuration, where upstream signals can occupy the 5 to 85 MHz band and downstream signals can occupy the 108 to 1218 MHz band. This will increase the upper frequency limit of the upstream signal path from the Node + N sub-split 42 MHz limit, to 85 MHz. This represents more than a 2x (43 MHz) increase, and an increase in capacity by more than 2x when CPE based on DOCSIS 3.1 1024-QAM OFDMA replacing CPE based on DOCSIS 3.0 64-QAM SC-QAM in the additional 43 MHz bandwidth, for over 500 Mbps total throughput.

The next upgrade transition will be to high-split networks, with upstream signals occupying the 5 to 204 MHz band. That yields an additional 119 MHz of upstream bandwidth, for over 1.5 Gbps total throughput. The downstream signals can occupy the 258 to 1218 MHz band. These high-split upgrades will move the video channels above 258 MHz prior to enabling high-split operation. This will prevent the spurious emissions interference problem previously discussed, as the video channels will be moved above the upstream cutoff frequency of the CM's diplex filter.

However, the legacy STBs will still have a sub-split diplex filter instead of the upgraded CM high-split diplex filter. Thus the cutoff frequency (54 MHz) of the sub-split STB's diplex filter will see signals in the 54 to 204 MHz upstream band from the high-split CMs. These signals will be adjacent to the 258 MHz-and-above STB video channels. If the CM's transmit power adjacent to the downstream video channels is strong enough, this can cause adjacent channel interference (ACI) which could impair the STB-delivered video channels.

This ACI scenario will also occur in the transition to mid-split operation previously described, but in the 54 to 85 MHz upstream band from the mid-split CMs. These signals will be adjacent to the 108 MHz-and-above STB video channels. The ACI power will be lower for the same power spectral density as the upstream bandwidth is limited to 85 MHz in the mid-split case. The total ACI power will be less than one-third of the high-split ACI power. Thus neighbor video interference is much less likely for mid-split operation.

If a customer desires a 1 Gbps upstream service with a high-split CM replacement, then any STBs present in that premises can be fitted with high-pass filters. The filters suppress the high transmit level of the high-split CM leaking across an in-home splitter, if needed. But, a neighbor on the same tap may see the increased ACI leaking across drops, as shown in Figure 10.



**Figure 10 – High-Split CM ACI into Neighbor Legacy STB**

**Table 5 – High-Split CM Carrier/ACI Ratio at STB Impaired Video Threshold**

Model	Measured Peak Power of Interferer (dBmV/96 MHz)	Measured Delta Peak to 6 MHz Video SC-QAM (dBc)	Corrected Power Delta (dBc)	Downstream Power (dBmV/6 MHz)
RNG150CNM	27.95	-34.5	-22.5	-5
DCT700	27.12	-35.79	-23.79	-5
DC50Xu	30.1	-36.27	-24.27	-5
SARNG100	26.35	-32.52	-20.52	-5
DC60Xu	29.17	-35.78	-23.78	-5
HD-DTA100u	28.36	-38.26	-26.26	-5
AX013ANM	30.03	-36.17	-24.17	-5
PX013ANM	30.01	-35.99	-23.99	-5
PX022ANM	28.98	-38.16	-26.16	-5
DCT2524	31.08	-37.85	-25.85	-5
DCX3200	25.09	-32.58	-20.58	-5
RNG110C	30.71	-37.82	-25.82	-5
RNG110RF	28.14	-35.62	-23.62	-5
DCI105COM	27.35	-36.08	-24.08	-5

The port-to-port isolation of the tap, plus the drop cable and in-home splitter insertion losses in Figure 10, determines the path loss between neighbors, which determines the attenuation of the CM-coupled ACI. The magnitude of the interference depends on each CM's transmit level, and the STB's receive level separated by this path loss. If the resulting ACI level from a CM exceeds the STB downstream receiver susceptibility threshold, then video artifacts or complete loss of reception of some channels above 258 MHz (that are adjacent to the CM's upstream transmit signal interference below 204 MHz) may occur for that customer. Measurements of the STB Carrier/ACI Ratio (CACIR) video threshold levels that exhibited "tiling" (macroblock freezing/errors in portions of the decompressed video) are shown in Table

5 for several deployed STB models at several frequencies. The video SC-QAM receive level was -5 dBmV/6 MHz for these measurements. The maximum CACIR STB threshold level where video artifacts are observed is below -20 dB across all models tested.

The same 23 million CM transmit and receive level statistics from the field are used for the analysis of the percentage of STBs with CM ACI resulting in CACIR below the -20 dB threshold of impaired video reception. The following analysis shows an estimated upper bound on the number of cases which have sufficient CM interference magnitude impairing STB video reception.

The data analysis method is as follows:

Calculate the average transmit and receive level per channel for each modem (per MAC address) over all upstream and downstream frequencies.

Estimate the isolation between neighbors (i.e. path loss  $P_{loss}$ ) as:

$$P_{loss} = \text{Tap Port-to-Port Isolation} + 2 * (\text{Drop Cable Attenuation} + \text{Splitter Insertion Loss})$$

Calculate the average ACI level for each average transmit level for each CM MAC address as:

$$\text{ACI} = \text{CM Transmit Level} - P_{loss}$$

Calculate the STB CACIR for each average downstream receive level and average upstream transmit level as:

$$\text{CACIR} = \text{CM Receive Level} - \text{ACI}$$

Sort the calculated STB CACIR levels in ascending order (plotted in Figure 11).

Calculate the percentage of CMs  $\leq$  each STB CACIR as:

$$\% \text{CMs} \leq \text{CACIR value} = \text{Number of CIR entries} \leq \text{CIR} / \text{Total Number of CIR values} * 100$$

**Table 6 – Analysis of Percentage of STBs Below Video Threshold CACIR**

CPE Parameters	
Splitter Insertion Loss (dB):	3.5
Splitter Isolation (dB):	35
Set-Top Box Threshold CNR (dB):	28
Set-Top Box Threshold CACIR (dB):	-20


Drop Coax Span (ft):	100
Drop Attenuation per 100 ft (dB):	1.33
Tap Port-to-Port Isolation (dB):	24

Threshold Video CACIR	
Total Downstream Band	
% Set-Top Boxes with < -20 dB CACIR	9.49%
Set-Top Box CACIR(dB):	-20.0

The cases which have excessive adjacent interference will require mitigation. The percentage of STBs below the -20 dB CACIR threshold for the specifications in Table 6 is estimated to be below 9.5%. These STBs would need high-pass filters on each tap port to block the ACI from neighboring CMs on the same tap.

However this is a conservative estimate as the higher CM upstream levels in many cases is due to higher attenuation in the path from the CM to the tap port including the in-home network which lowers the interference level into the neighboring tap ports.






**Figure 11 – Percentage of STBs ≤ CACIR**

## 6. Accommodating Legacy MoCA in 1.2 GHz System Upgrades

Legacy STBs have an SC-QAM upper frequency limit of 1002 MHz. In these legacy designs, the MoCA channels are situated starting with a lower band edge frequency of 1125 MHz. Mid-split and high-split systems with DOCSIS 3.1 CMs, as well as forthcoming DOCSIS 4.0 FDX CMs, can support an extended OFDM upper frequency limit above 1002 MHz to 1218 MHz. In order to utilize this additional DOCSIS bandwidth, MoCA channels have to migrate to higher channel frequencies.

Figure 12 illustrates the lowest non-overlapping MoCA channel and the DOCSIS 3.1 OFDM upper frequency limit. The minimum separation for non-overlapping MoCA channels is from the edge of the MoCA channel to the edge of the second adjacent emissions sideband equal to 75 MHz. The DOCSIS 3.1 highest maximum bandwidth (192 MHz) OFDM channel extends from the DOCSIS 3.0 upper limit of 1002 MHz to 1194 MHz (not up to the DOCSIS 3.1 1218 MHz upper limit).

The MoCA Extended Band D frequency plan is shown in Table 7. The lowest frequency MoCA 1.1 and 2.0 channels with a bandwidth of 50 MHz and 100 MHz respectively, that are separated by at least the minimum non-overlapping channel gap of 75 MHz, are noted in the table. The minimum gap from the respective DOCSIS and MoCA channel boundaries is 81 MHz.





**Figure 12 – MoCA Non-Overlapping Channel & DOCSIS 3.1 Downstream Frequency Limit**

The MoCA Alliance performed a MoCA 2.0 field trial in a number of homes to determine the percentage of coax outlets that can support a given MAC rate over all the D band channels [3]. The results for Comcast participants extracted from that field trial report data is shown in Figure 13. The D band channels D4 through D6 in the 1275 MHz to 1475 MHz frequency band were shown to deliver a 400 Mbps MAC rate in over 90% of measured outlets.

In order to place the DOCSIS and MoCA channels in such close proximity, the RF isolation requirements and the resulting trade-offs in DOCSIS 3.1 SNR are analyzed in the following. An evaluation of required isolation between MoCA transmit and DOCSIS 3.1 receive signals in the common RF port is performed for two sources of adjacent channel interference: Susceptibility (maximum power of signals adjacent to the desired signal) and Noise and Spurious Emissions (out-of-band interference from signal emissions adjacent to the desired signal). The DOCSIS 3.1 PHY spec and the MoCA 2.0 spec determine the signal and emissions levels and bandwidth tabulated in Table 8 and Table 9 respectively.

**Table 7 – MoCA Extended Band D Frequency Plan**



**Figure 13 – MoCA 2.0 Field Trial Results of Percentage of Outlets vs. MAC Rate**

## Table 8 – DOCSIS 3.1 RF Specifications

*Table 1 – DOCSIS 3.1 Downstream Signal Requirements*

Parameter	Value
Lower Frequency Boundary	256 MHz Required 108 MHz Optional
Upper Frequency Boundary	1218 MHz Required 1794 MHz Optional
CM's Minimum Receive Power	-15 dBmV/6 MHz
CM's Maximum Receive Power	15 dBmV/6 MHz
Maximum average power per MHz input to the CM (dBmV/MHz)	Min $[X - 10 \log(24) + 10; 21 - 10 \log(24)]$ where X = Average power of lowest power 24 MHz bandwidth for demodulation; $X \leq 11$ dBmV.
Maximum out of Channel Emissions for DOCSIS 3.1 192 MHz Channel (from DOCSIS spec with N=32):	
0 – 750 kHz from Channel Edge	-57.4 dBr
750 kHz – 6 MHz from Channel edge	-56.5 dBr
6 -12 MHz from Channel Edge	-57.3 dBr
12 -18 MHz from Channel Edge	-57.9 dBr
> 18 MHz from Channel Edge Note: dBr ≡ dB relative to the DOCSIS signal power in a 6 MHz bandwidth	-57.9 dBr
Maximum out of Channel Emissions for DOCSIS 3.1 192 MHz Channel into MoCA Channels (calculated from the values above for the maximum received signal level):	
5– 55 MHz from Channel Edge	-33.1 dBmV
5 – 105 MHz from Channel Edge	-30.4 dBmV
25 – 75 MHz from Channel Edge	-33.7 dBmV
75 – 175 MHz from Channel Edge	-30.7 dBmV

*Table 7-41 - CM Minimum CNR Performance in AWGN Channel*

Constellation	CNR <sup>1,2</sup> (dB) Up to 1 GHz	CNR <sup>1,2</sup> (dB) 1 GHz to 1.2 GHz	Min P <sub>avg</sub> dBmV
4096	41	41.5	-6
2048	37	37.5	-9
1024	34	34	-12
512	30.5	30.5	-12
256	27	27	-15
128	24	24	-15
64	21	21	-15
16	15	15	-15

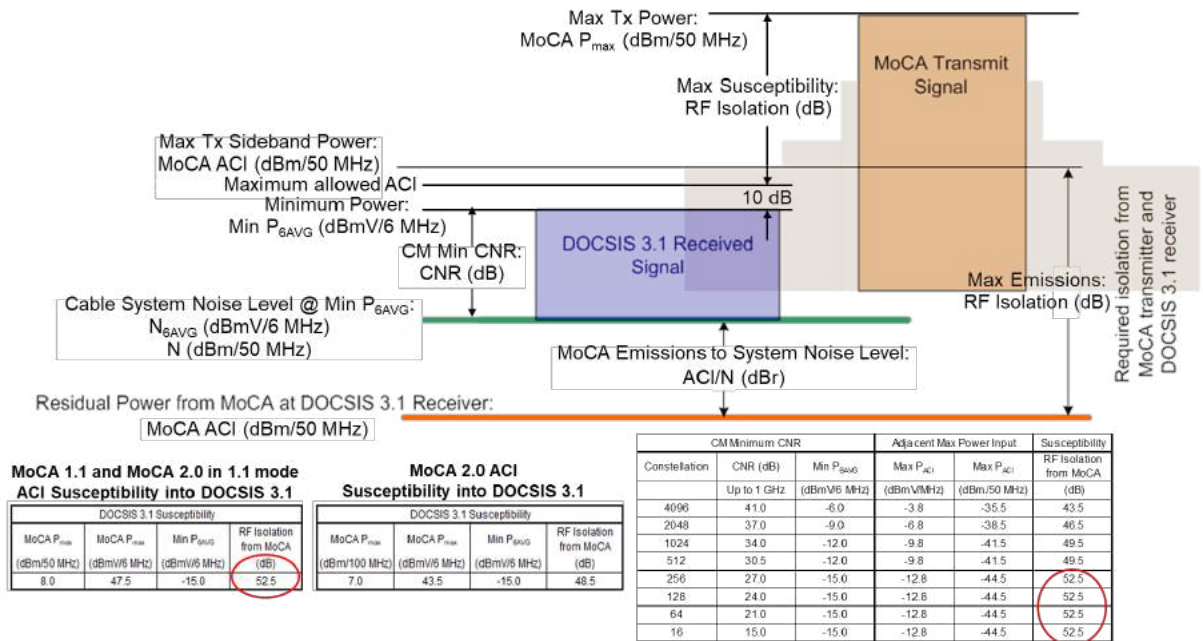
Table Notes:  
 Note 1 CNR is defined here as total signal power in occupied bandwidth divided by total noise in occupied bandwidth  
 Note 2 Channel CNR is adjusted to the required level by measuring the source inband noise including phase noise component and adding the required delta noise from an external AWGN generator

## Table 9 – MoCA RF Specifications

*Table 2 – MoCA Signal Specifications*

Parameter	Value
MoCA 1.1 Channel Width	50 MHz
MoCA 1.1 Device Maximum Transmit Power	8 dBm/50 MHz
MoCA 2.0 Channel Width	100 MHz
MoCA 2.0 Device Maximum Transmit Power	7 dBm/100 MHz 10 dBm/225 MHz (2 x 100 MHz channels separated by 25 MHz)
Band ExD Lower Frequency Boundary	1125 MHz
Band ExD Upper Frequency Boundary	1675 MHz
Out of Channel Emissions	Adjacent channels: MoCA 1.1 channels: -40dBr MoCA 2.0 channels: -40dBr Non overlapping frequencies: MoCA 1.1 channels: -45dBr MoCA 2.0 channels: -50dBr
Receiver ACI Susceptibility	First Adjacent Channel : Interference power < S - 6 dB  Second Adjacent Channel: Interference power < MAX{ - 42 dBm, MIN[S+8, -2 dBm]}; Where S is the received MoCA signal level
Minimum Receive Level	MoCA 1.1 channels: ~ -67 dBm MoCA 2.0 channels: ~ -64 dBm

MoCA adjacent channel susceptibility into DOCSIS 3.1 analysis and results follow the methodology in [2] and are summarized in Figure 14. The MoCA transmitter to DOCSIS 3.1 receiver minimum RF isolation to meet the limits of susceptibility (ACI) between these signals is 52.5 dB.



**Figure 14 – Analysis of MoCA Susceptibility into DOCSIS 3.1**

The other RF isolation requirement is from noise and spurious emissions interference from a MoCA 2.0 transmitter, into a DOCSIS 3.1 receiver. Our analysis and results of MoCA 2.0 and MoCA 2.0 in 1.1 mode for non-overlapping channel noise and spurious emissions into DOCSIS 3.1 follow the methodology in [2]. The methodology evaluates the MoCA emissions impact on DOCSIS 3.1:

- A maximum tolerable level of degradation to the input CNR is chosen  $\Delta\text{CNR} = [C/N]/[C/(N+ACI)]$  dB corresponding to a maximum ACI emissions-to-noise ratio ACI/N
- The input noise floor is calculated for any spectral efficiency (modulation order) at its minimum required DOCSIS 3.1 signal level
- The required RF isolation is calculated over all modulation orders for this maximum input CNR degradation or less at the calculated ACI/N ratio and maximum MoCA ACI emissions level

The required RF isolation, so calculated, should be greater than or equal to the highest RF isolation required for all other ACI susceptibility cases previously considered (i.e. 52.5 dB MoCA ACI susceptibility into DOCSIS 3.1.)

The analysis and results are summarized for MoCA 2.0 and MoCA 2.0 in 1.1 mode in Figure 15 and Figure 16, respectively. The MoCA transmitter to DOCSIS 3.1 receiver minimum RF isolation to meet the limits of noise and spurious emissions into DOCSIS 3.1 for non-overlapping channels between these signals is 52.5 dB. The  $\Delta\text{CNR}$  limits to match the previous susceptibility isolation requirement are 0.25 dB for the MoCA 2.0 transmitter to DOCSIS 3.1, and 0.65 dB for the MoCA 2.0 in 1.1 mode transmitter to DOCSIS 3.1.



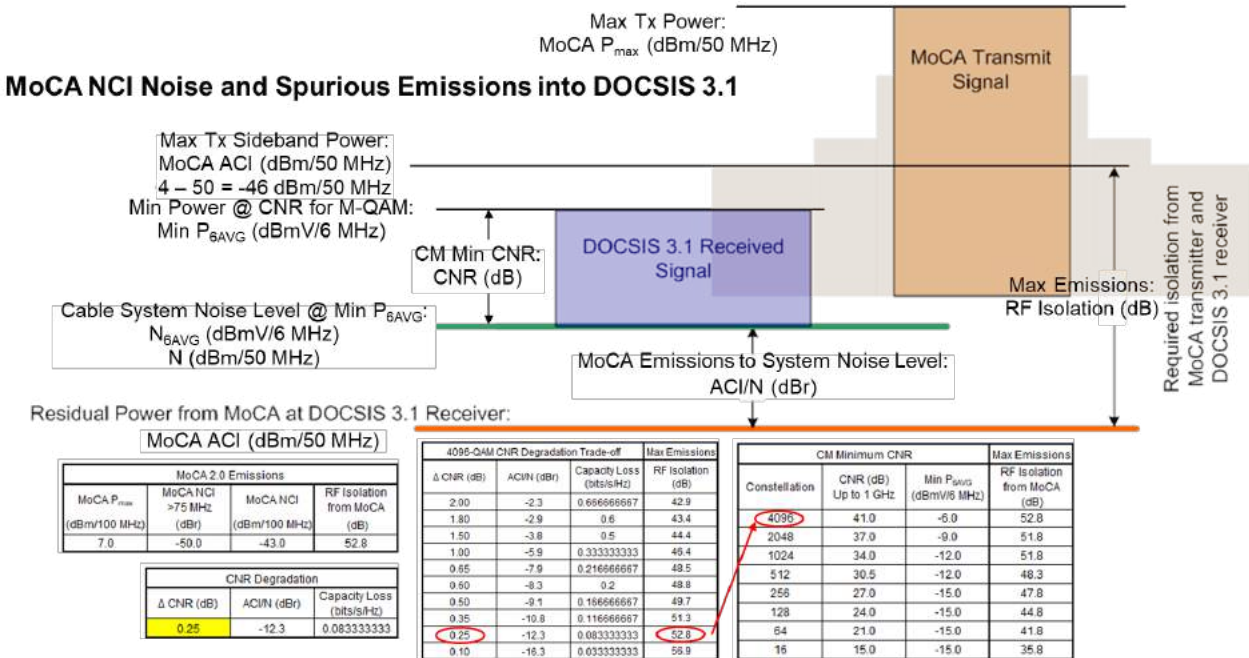


Figure 15 – MoCA 2.0 Transmitter to DOCSIS 3.1 Receiver RF Isolation

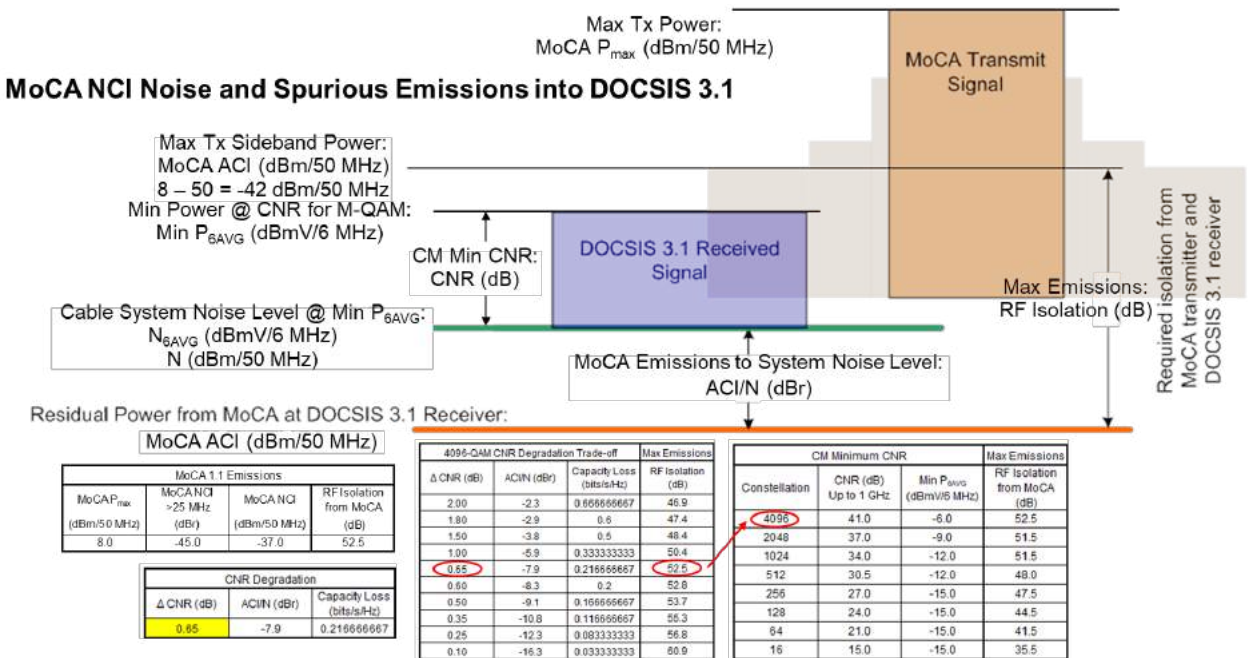


Figure 16 – MoCA 2.0 in 1.1 mode Transmitter to DOCSIS 3.1 Receiver RF Isolation

Incorporating the CNR requirements for DOCSIS 3.1 provides the required RF isolation between a MoCA transmitter and a DOCSIS 3.1 receiver in non-overlapping channels, resulting in a practical isolation filter with minimal loss in spectral efficiency (bit-loading or bits/s/Hz), or alternatively maintaining spectral

efficiency with an increased minimum DOCSIS receive level where increased minimum receive level = nominal minimum receive level + CNR loss ( $\Delta\text{CNR}$ ).

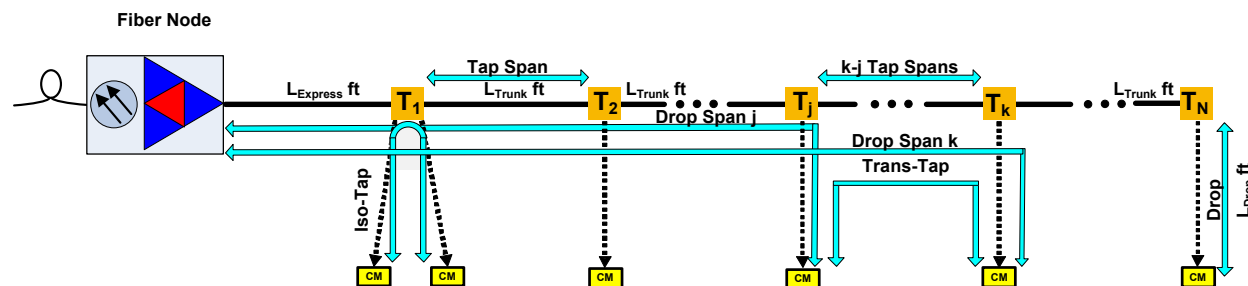
MoCA Noise and Spurious Emissions into DOCSIS 3.1 for minimally impacting the required CNR at the minimum receive level of the highest spectral efficiency (modulation order) in the DOCSIS 3.1 specification requires a minimum RF isolation equal to the 52.5 dB minimum susceptibility interference isolation. This requirement is achieved with 0.25 to 0.65 dB CNR loss for Non-Overlapping Channel (NCI) with > 75 MHz transition band with MoCA 2.0. This can be implemented with a DOCSIS 3.1/MoCA triplex filter incorporated into the in-home network CPE, with at least 52.5 dB of stop band attenuation in the DOCSIS 3.1 band pass and MoCA high pass filter sections, with an 81 MHz transition band and a low pass legacy upstream band filter section.

## 7. Upgrading to Full Duplex DOCSIS for the Node + 0 Architecture

The transition from either mid-split or high-split networks to FDX, with a progression of increasing FDX bandwidth from 192 MHz to 576 MHz, is described in this section.

The DOCSIS 4.0 FDX architecture provides two-way signal transmission within the same spectral band. This requires a passive architecture without amplifiers. In this case, the fiber node connects to a single series of multiport taps. Without any amplifiers that require duplex filtering, both upstream and downstream signals can share the part of the same spectrum (108 MHz to 684 MHz in up to three 192 MHz sub-bands in the FDX band) but with the same directivity of the conventional architecture.

The FDX DOCSIS PHY layer design and performance detailed analysis is discussed in [1]. An example of one coaxial branch of a Node + 0 architecture and model used in [1] is shown in Figure 17.



**Figure 17 – FDX Interference Analysis Model in the Node + 0 Network**


This model is used to analyze co-channel interference (CCI) between cable modems (CMs) both across the same tap and different taps to calculate SNR and bit-loading (i.e. spectral efficiency). Also calculated is the self-interference resulting from transmitting and receiving signals including reflections (echoes) within the same frequency bands in the node, and adjacent channel and leakage interference (ACI and ALI) within adjacent frequency bands in each modem. It is shown that these interference levels need to be sufficiently suppressed through Echo Cancellation (EC) to obtain high SNR and bit-loading of both the upstream and downstream signals in the FDX band.

The DOCSIS 4.0 FDX allocated spectrum bandwidths can be configured as shown in Figure 18. Non-allocated spectrum is shown within the FDX band between the highest frequency of the FDX allocated spectrum and 684 MHz. There are no US channels in the non-allocated spectrum of the FDX band.

An FDX Node must be able to place either SC-QAM or non-FDX OFDM channels in the non-allocated spectrum of the FDX band with RF performance requirements specified in this section.

An FDX CM in FDX mode is not required to receive either SC-QAM or non-FDX OFDM channels in the non-allocated spectrum of the FDX band.

The FDX transition band is the spectrum between 684 MHz and 804 MHz. It is possible that performance in this band could be impaired compared to a compliant DOCSIS 3.1 CM. As such, performance in the FDX transition band for an FDX CM operating in FDX mode is not specified.

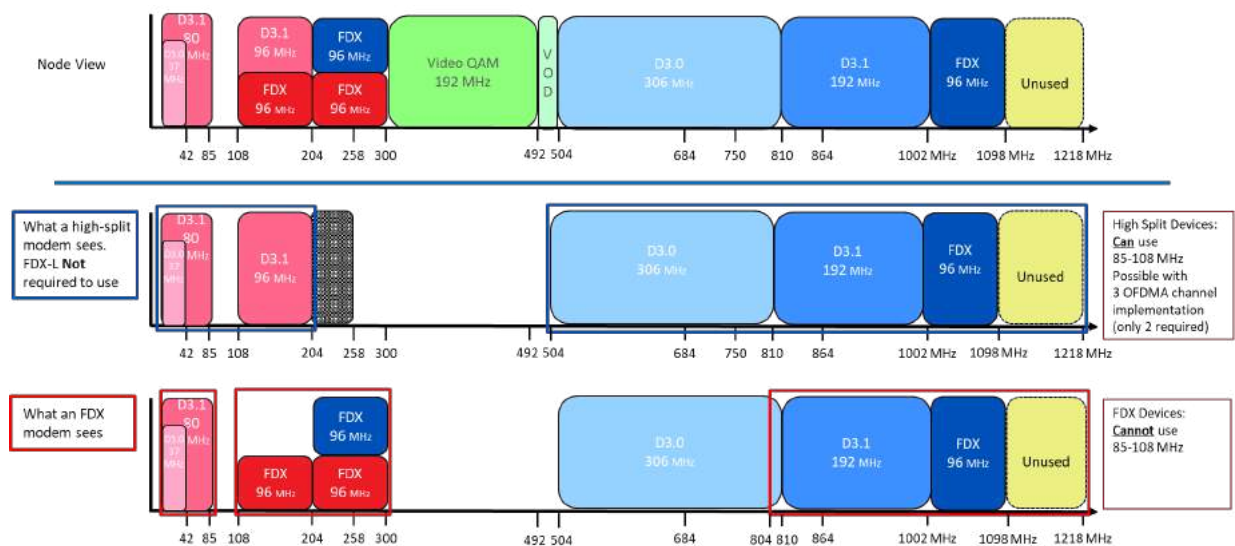


**Figure 18 – Configurable FDX Allocated Spectrum Bandwidths**

An example of an initial FDX spectrum allocation for the 192 MHz allocated spectrum of Figure 18 is shown in Figure 19. The DOCSIS 4.0 FDX specification always starts the first of three sub-bands in the FDX band at 108 MHz. Video SC-QAM channels are supported in the lowest downstream frequency band above the FDX allocated spectrum bandwidth. This initial upgrade, supporting one allocated FDX sub-band, has two Resource Blocks (RBs). Each can be set to either the upstream or downstream direction.

The first Resource Block Assignment (RBA) in the 108 MHz to 204 MHz is set to a static upstream direction. The second RBA in the 204 MHz to 300 MHz sub-band can be set dynamically to either the downstream or upstream direction. Full duplex operation in the network branch of Figure 17 is therefore possible within this second RB, if two or more Interference Groups (IGs) can be divided up between two Transmission Groups (TGs) that have low Co-Channel Interference (CCI). This would allow one CM in one TG to transmit upstream while all CMs in the other TG receive downstream.





**Figure 19 – RF Spectrum: FDX with Sub-Split, Mid-Split and High-Split – 1.1 GHz Plant**

The reason to set the first RB to a static upstream direction is to allow high-split CMs that were added to the network before the initial FDX upgrade to continue to share the 108 MHz to 204 MHz band with FDX CMs. If this RB could be assigned to the downstream direction and if the high-split CM was in the same TG, then that CM could not transmit when the RBA was set to the downstream direction.

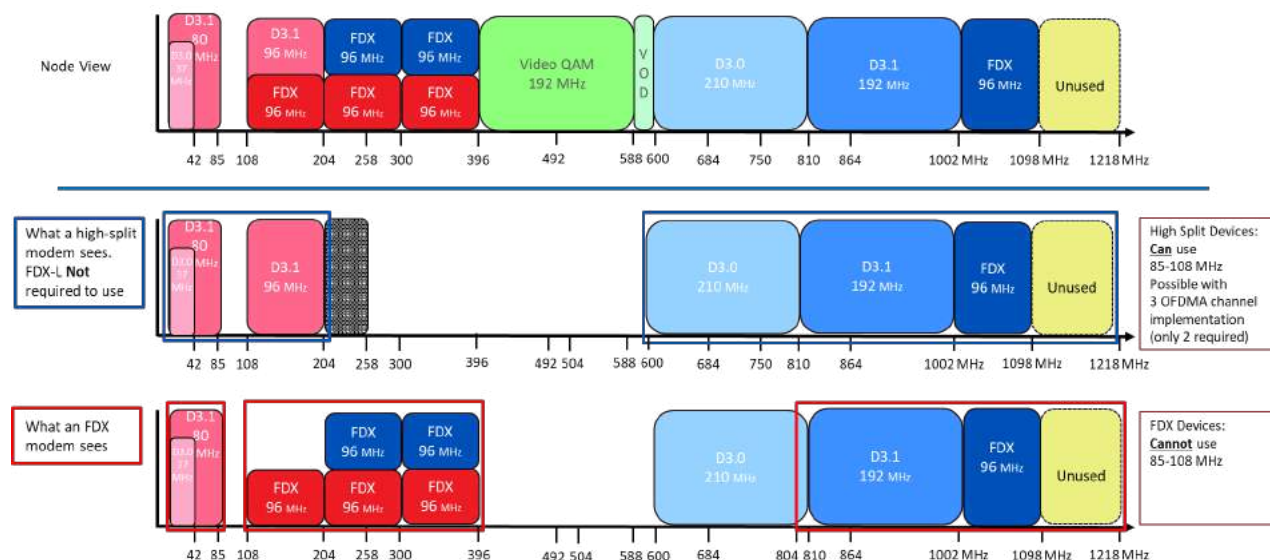
In order to determine if a high-split CM was in the same TG as an FDX CM, the high-split CM would have to participate in the sounding process, along with the FDX CMs. This sounding capability in a high-split CM requires a software upgrade to a legacy DOCSIS 3.1 cable modem (FDX-L) which can transmit in the 108 MHz to 204 MHz Full Duplex upstream channels and receive in the 258 MHz to 684 MHz Full Duplex downstream channels in a high-split access network, or it can receive in the 108 to 684 MHz Full Duplex downstream channels in a mid-split access network with no access to upstream Full Duplex Channels.

As shown in Figure 19, FDX downstream reception is disallowed in the first RB. Only upstream transmissions from either high-split CMs or FDX CMs is allowed. High-split downstream reception is only allowed above the allocated FDX band, where no upstream transmission are present in the DOCSIS 3.0 SC-QAM band above 504 MHz. Therefore no FDX-L software upgrade is needed to support sounding in the high-split CMs. There is never an opportunity for a high-split CM to transmit in a frequency that can be used to receive downstream by FDX CMs. Conversely there is no FDX upstream transmission in the high-split downstream band above the FDX allocated spectrum band.

As previously stated above, an FDX CM in FDX mode is not required to receive either SC-QAM or non-FDX OFDM channels in the non-allocated spectrum of the FDX band. So an FDX CM operating in the RF spectrum allocation of Figure 19 receives downstream above the FDX transition band starting at 804 MHz.

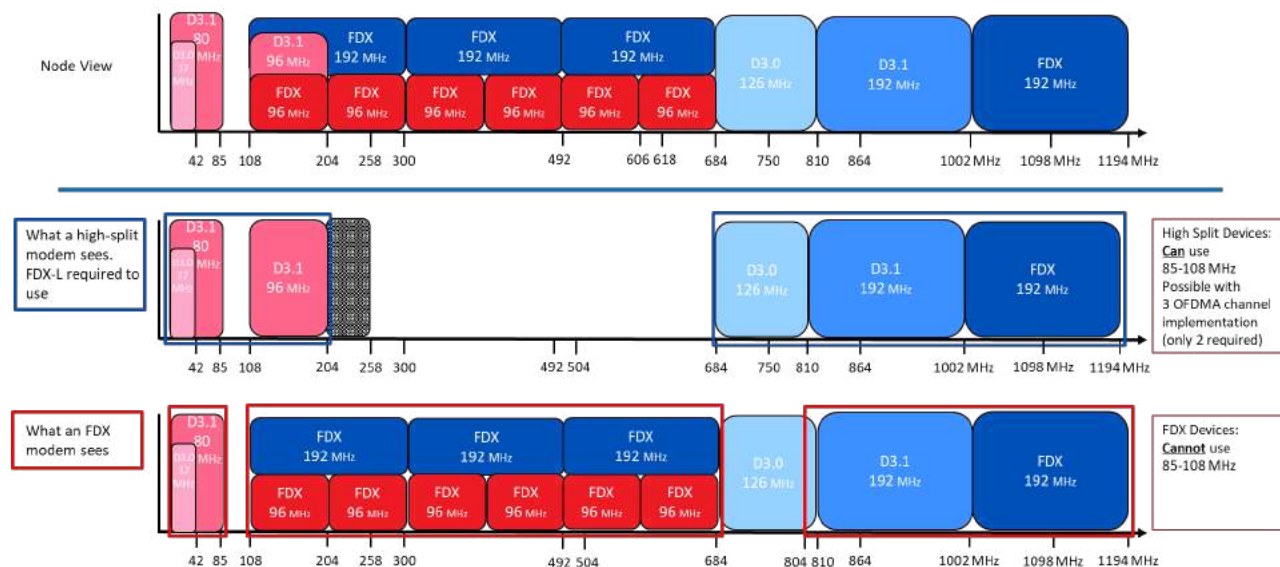
This initial FDX upgrade to a Node + 0 node provides FDX CMs with 192 MHz of additional upstream spectrum to the 5 MHz to 85 MHz legacy upstream spectrum and 192 MHz of additional FDX downstream spectrum to the legacy high-split downstream spectrum. This upgrade also provides high-split CMs with 96 MHz of additional FDX downstream spectrum to the legacy high-split downstream

spectrum -- but loses 23 MHz of the upstream spectrum from 85 MHz to 108 MHz, in the original 199 MHz upstream band from 5 MHz to 204 MHz.



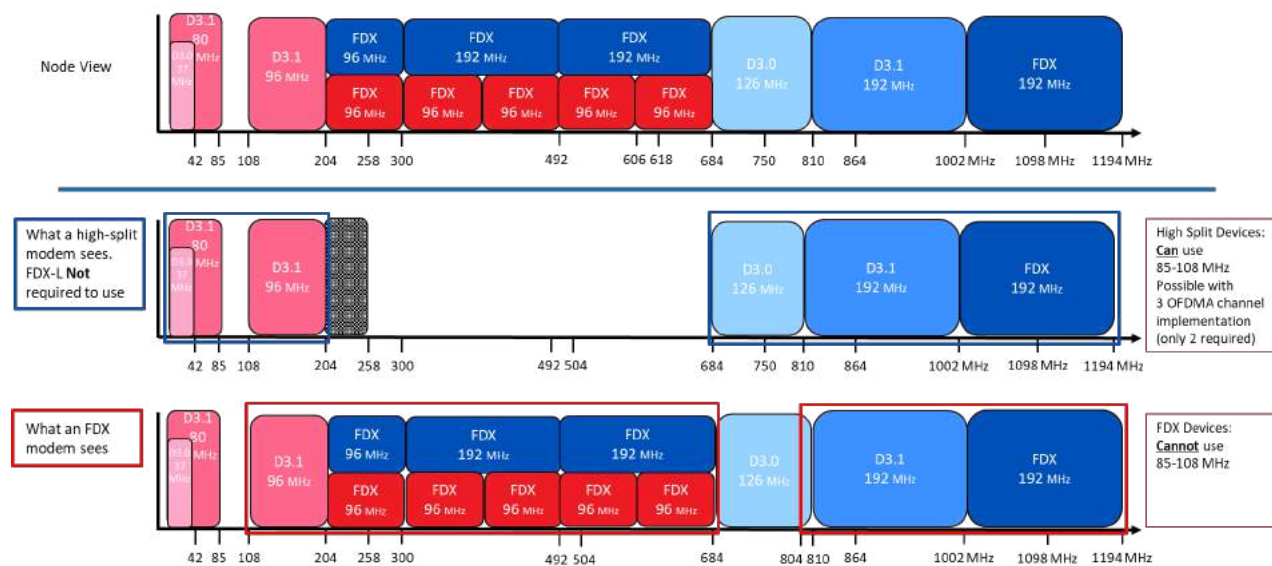
**Figure 20 – RF Spectrum: FDX with Sub-Split, Mid-Split and High-Split – 1.2 GHz Plant**

Future demand for bandwidth growth can be provided by reallocating the same total RF spectrum as shown in Figure 20. This next upgrade preserves the video SC-QAM channels moved up by 96 MHz to allocate that spectrum to a third 96 MHz FDX RB. The DOCSIS 3.0 SC-QAM channels are reduced by the same 96 MHz as more customers choose to upgrade their DOCSIS 3.0 service tier to the high-split tier or higher still to the FDX tier.



**Figure 21 – RF Spectrum: Full FDX with Sub-Split, Mid-Split and High-Split – 1.2 GHz Plant (Compliant with DOCSIS 4.0 Spec)**

Continued demand for more bandwidth in the future can be provided by reallocating the same total RF spectrum as shown in Figure 21. This next upgrade removes the video SC-QAM channels to allocate that spectrum to a third 192 MHz FDX RB, and doubles the bandwidth of each of the first two FDX RBs to 192 MHz. The DOCSIS 3.0 SC-QAM channels are reduced by 84 MHz as more customers choose to upgrade their DOCSIS 3.0 service tier to the high-split tier, or, higher still, to the FDX tier. However, in this upgrade, the high-split CMs must be software-upgraded to FDX-L, because the 108 MHz to 204 MHz high-split upstream shares some of the downstream FDX bandwidth in the first FDX RB.



**Figure 22 – RF Spectrum: Full FDX with Sub-Split, Mid-Split and High-Split – 1.2 GHz Plant (Non-Compliant with DOCSIS 4.0 Spec)**

An alternative to the full FDX upgrade of Figure 21 is shown in the nearly full FDX upgrade of Figure 22. In this RF configuration, the FDX band starts at 204 MHz instead of 108 MHz, sacrificing 96 MHz of FDX bandwidth. This obviates the need to upgrade the software of all high-split CMs to FDX-L, as the high-split CMs do not share any upstream bandwidth with the FDX downstream band. However, this configuration does not start the FDX lower band edge, at 108 MHz, which, strictly speaking, is not compliant with the DOCSIS 4.0 spec. But this should not require a hardware change to the FDX CM or CMTS.

## 8. Optimizing Full Duplex Architecture and Operation

### 8.1. FDX Modem Downstream Legacy Band Receive Power

In December 2018, a proposal to the CableLabs FDX Working Group by CM silicon providers with two options, A and B, was presented for the CM architecture, which proposed changing either the CM input minimum return loss or the maximum insertion loss into the CM F-connector when operating in FDX mode. The two options discussed are shown in Figure 23. The difference in the options was in the placement of the directional coupler for splitting and combining upstream and downstream signals in the FDX band.

If the FDX diplex filter common port is connected to the CM F-connector, and the diplex filter legacy downstream port is connected to the legacy downstream receiver, as shown in Option A, then the only insertion loss seen at the legacy band receiver is the small diplex filter insertion loss. However, the return



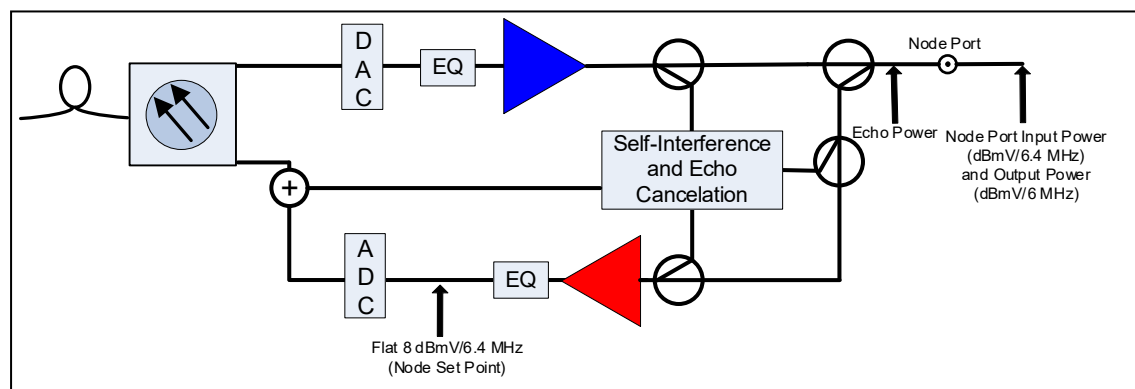
Multiple vendors obtained filters from third parties that meet the 6 dB minimum return loss spec across the entire band from 108 MHz to 1218 MHz. Thus the DOCSIS 4.0 legacy downstream spec was restored to the original DOCSIS 3.1 spec in the 804 MHz to 1218 MHz band.

## 8.2. FDX Modem FDX Band Receive Power

The directional coupler tap insertion loss for the legacy downstream can be avoided with the CM architecture of Figure 24 discussed in the previous section. However, the FDX downstream band will necessarily incur the insertion loss of the directional coupler tap. There are several mitigation approaches to increase the input power level to the FDX downstream receiver:

1. Decrease the directional coupler tap loss with a splitter or power divider/combiner
2. Switch the coupler inputs so the through port feeds the downstream receiver and the power amp transmits through the tap port
3. Increase the FDX band downstream power spectral density to counteract the coupler tap loss.

The first two approaches have the drawback that they require significantly higher total composite power from the upstream power amplifier. This total power spread over a wide bandwidth up to  $684 - 108 = 576$  MHz is specified to supply 65 dBmV minimum. Increasing power amp output by 4 to 10 dB significantly increases the modem input power dissipation and heat. The last approach does not need to modify the modem architecture and places the increased downstream power requirements on the node.



**Figure 25 – FDX Remote PHY Node Functional Block Diagram**

Figure 25 depicts a high level functional architecture for an FDX Remote-PHY node. To overcome the frequency-dependent attenuation of the hardline coax, which increases with increasing frequency, the downstream DAC provides digitally-generated up tilt and may be cascaded with an analog equalization filter (EQ) at the input of the downstream power amplifier (PA).


The resultant power spectral density (PSD) for two PSD vs frequency profiles at the node port is illustrated in Figure 26 and Figure 27. The nominal FDX PSD profile for a Node + 0 node output is a linear straight line, with 21 dB up tilt from 37 dBmV/6 MHz at 108 MHz, to 58 dBmV/6 MHz at 1218 MHz, with a Total Composite Power of 73.8 dBmV in both figures.

Figure 26 shows a “sawtooth” shape with a 7 dB boosted level over the nominal FDX PSD profile in the FDX band, from 108 MHz to 684 MHz, and a 3 dB attenuated level below the nominal FDX PSD profile from 876 MHz to 1218 MHz. The 7 dB boost matches the 7 dB attenuation of the CM directional coupler tap port in the FDX band to restore the receive level equal to the legacy downstream band, where legacy



DOCSIS 3.0 and 3.1 CMs receive in the 684 MHz to 876 MHz band for 750 MHz and 860 MHz cable plants.

Note that the level for legacy modems is unchanged in the band from 684 to 876 MHz where modems receive their downstream currently. Also note that the Total Composite Power of 73.8 dBmV is preserved with a larger boost at the lower FDX frequencies and a smaller attenuation at the highest frequencies above the current legacy downstream band. The reason for this behavior is that most of the power loading results at the highest frequencies due to the uptilted PSD. So no additional output power from the amplifier is required when balancing the higher boosted low frequencies with the lower attenuated high frequencies.




Downstream Start Frequency (MHz):	108
Downstream Stop Frequency (MHz):	1218

Power @Downstream Start Freq 108 MHz:	44.0
Tilt Step-Down Low Frequency (MHz):	684
Tilt Step-Down Low Amplitude (dB):	7.0
Tilt Step-Down High Frequency (MHz):	876.0
Tilt Step-Down High Amplitude (dB):	10.0

Total Composite Power (dBmV) - Linear:	80.8
Total Composite Power (dBmV) - Sawtooth:	73.7
Total Composite Power (dBmV) - FDX:	73.8

**Figure 26 – FDX Node RF Level with FDX Band Downstream Level Boosted +7 dB**



Downstream Start Frequency (MHz):	108
Downstream Stop Frequency (MHz):	1218

Power @Downstream Start Freq 108 MHz:	41.0
Tilt Step-Down Low Frequency (MHz):	684
Tilt Step-Down Low Amplitude (dB):	4.0
Tilt Step-Down High Frequency (MHz):	1002.0
Tilt Step-Down High Amplitude (dB):	5.0

Total Composite Power (dBmV) - Linear:	77.8
Total Composite Power (dBmV) - Sawtooth:	73.8
Total Composite Power (dBmV) - FDX:	73.8

**Figure 27 – FDX Node RF Level with FDX Band Downstream Level Boosted +4 dB**

Figure 27 shows a “sawtooth” shape with a 4 dB boosted level over the nominal FDX PSD profile in the FDX band, from 108 MHz to 684 MHz, and a 1 dB attenuated level below the nominal FDX PSD profile from 1002 MHz to 1218 MHz. The 4 dB boost partially restores the 7 dB attenuation of the CM

directional coupler tap port in the FDX band, to partially restore the receive level closer to the legacy downstream band where legacy DOCSIS 3.0 and 3.1 CMs receive in the 684 MHz to 1002 MHz band for 1 GHz cable plants. The 3 dB reduction in level will reduce the received OFDM bit-loading by about 1 bit per symbol.

Note that the level for legacy modems is again unchanged in the band where modems receive their downstream currently. Also note that the Total Composite Power of 73.8 dBmV is preserved in this case as well with a larger but less than the previous boost at the lower FDX frequencies and a smaller but less than the previous attenuation at the highest frequencies above the current legacy downstream band in 1 GHz cable plant. So no additional output power from the amplifier is required in either case.

### 8.3. Signal-to-Echo Ratio Optimization

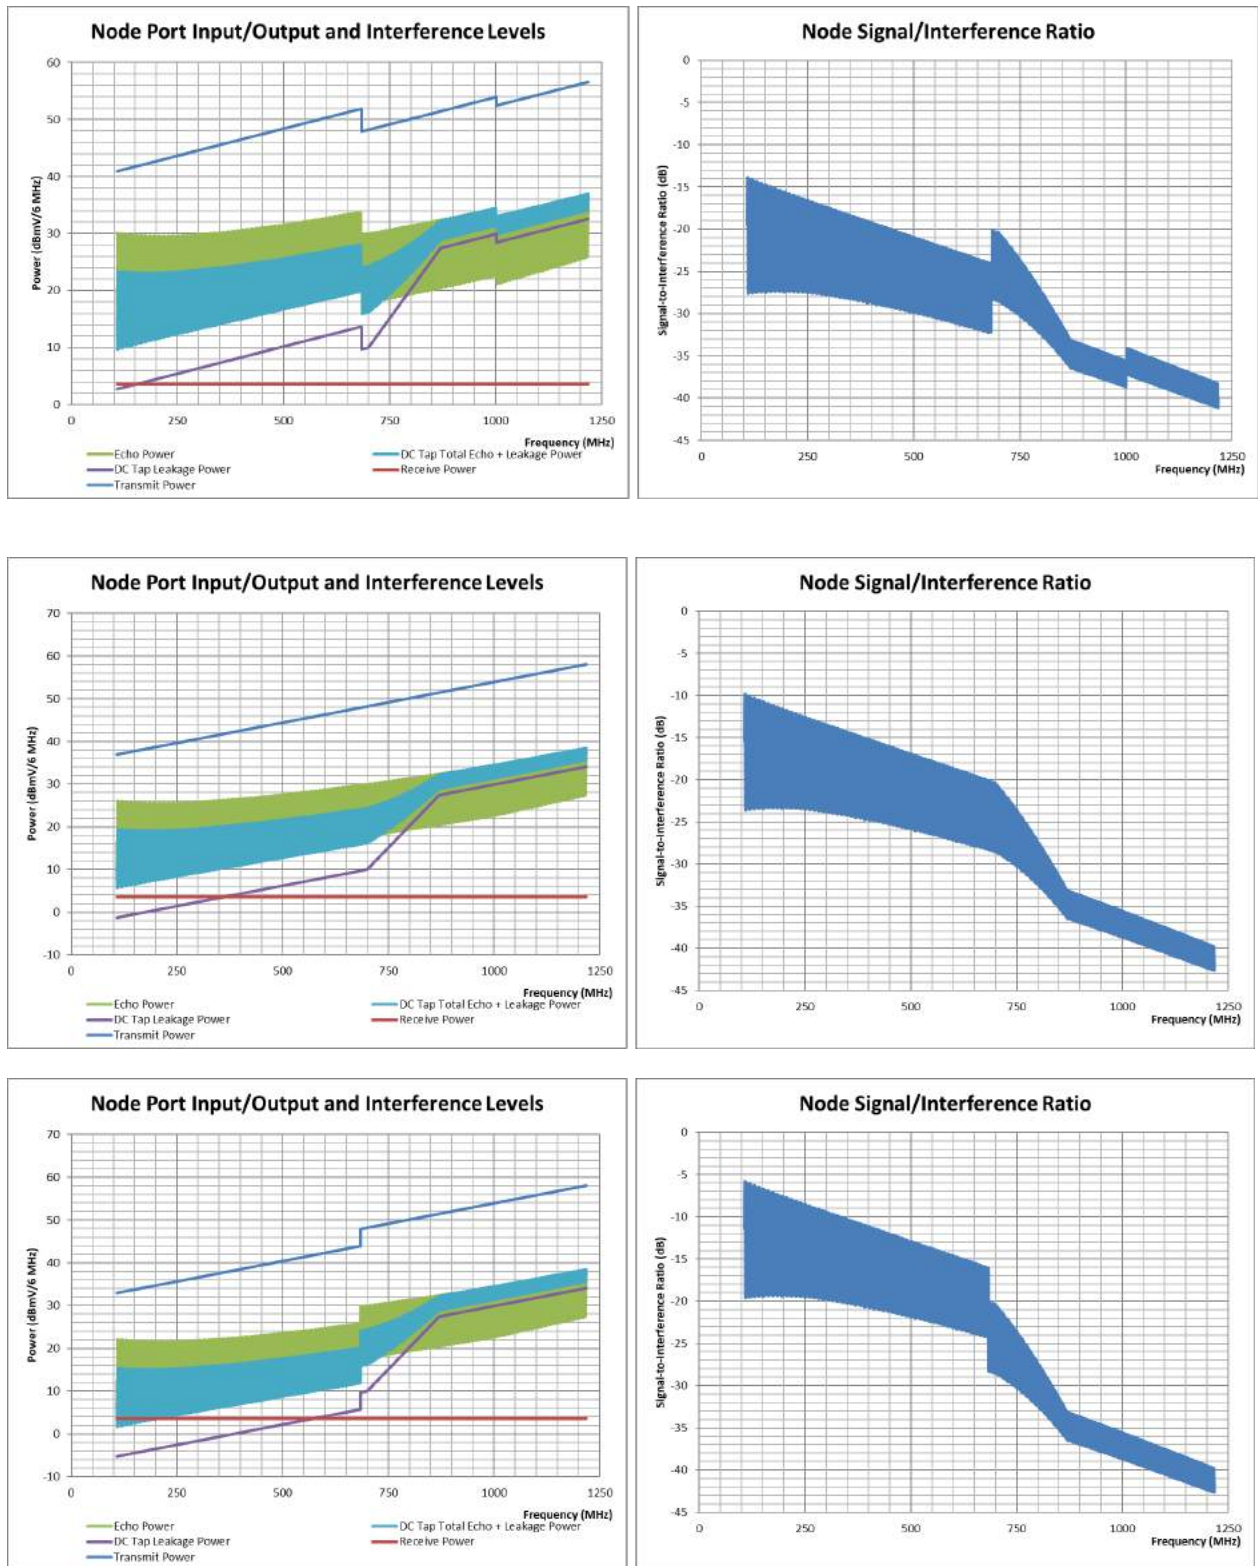
In the previous section, it was noted that increasing the node downstream PSD level increases the CM (downstream) Signal-to-(Upstream) Interference Ratio. The CM interference consists of self-ACI and its echoes (reflections) plus Adjacent Leakage Interference (ALI) and its echoes when transmitting upstream signals in one RBA while receiving downstream signals in another. The CM never transmits and receives in the same RBA whose direction is assigned either upstream or downstream dynamically. The CM essentially is working in frequency division duplex (FDD) mode but in a dynamically scheduled way without switched diplex filters and their associated transition bands but instead with echo cancellation and no guardbands and their associated spectral efficiency loss.

While this will increase the Signal-to-Echo Ratio and spectral efficiency (bit-loading) in the CM downstream receiver, the increase in the PSD in the node will decrease the Signal-to-Echo Ratio in the node upstream receiver. This ratio in the node is much lower as the upstream and downstream signals are full duplex with the co-channel interference ratio being orders of magnitude higher than non-overlapping adjacent signal emissions of the CM. Thus the node requires very high dynamic range echo cancellation.

Figure 28 shows an FDX network simulation we developed of the Upstream Node Signal/Interference Ratio vs FDX Downstream Band PSD Boost. The top pair of graphs are for a +4 dB boost; the middle pair of graphs are for no boost or a linear uptilt PSD; and the bottom pair of graphs are for a -4 dB attenuation. The left side shows all upstream and downstream signal levels, downstream echoes, and leakage across the FDX directional coupler ports, and the total combined interference level into the upstream received signal at the node. The right side shows the resulting calculated upstream Signal-to-Interference Ratio (SIR) of the node. These results show the direct correlation of the downstream node PSD level in the FDX band and the resulting upstream SIR of the node in that band.

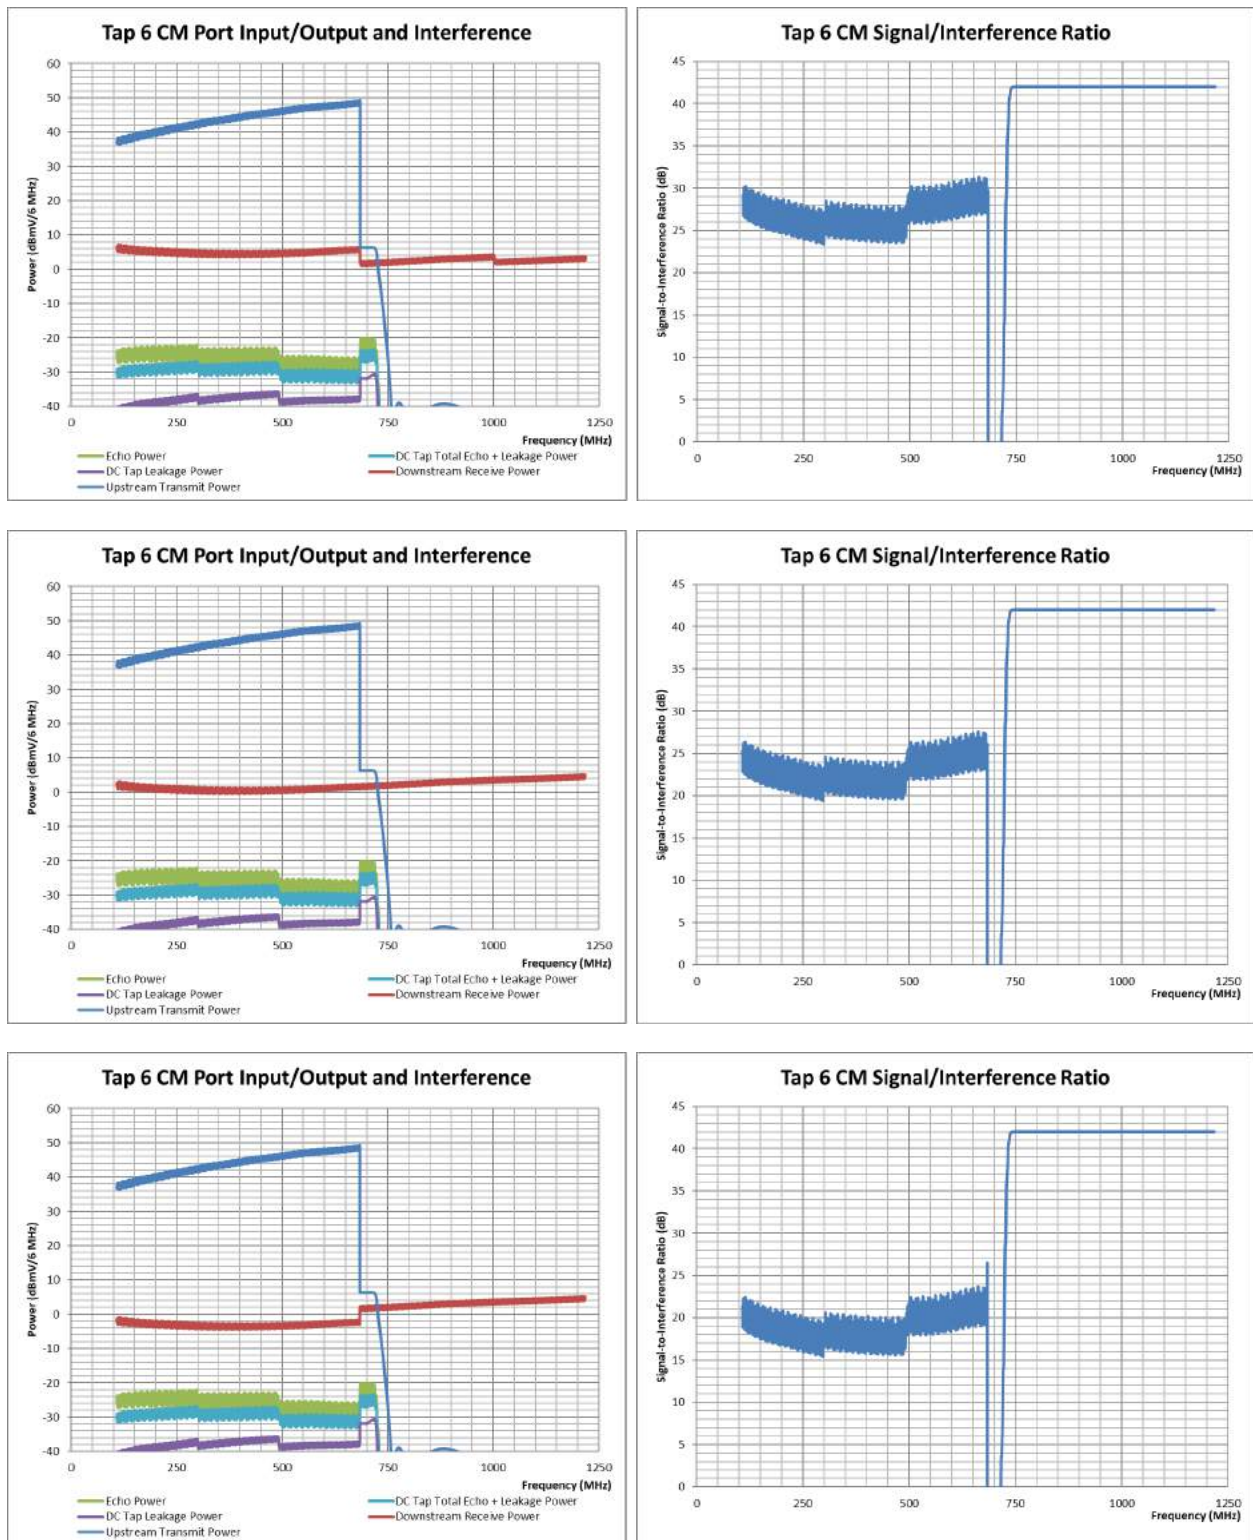
For a +4 dB downstream PSD boost, the SIR averages from -20 dB at 108 MHz to -28 dB at 684 MHz. The tilt in the SIR is due to the downstream PSD uptilt launched by the node power amp and relatively flat upstream received signal. For a 0 dB downstream PSD boost (strictly linear without step-downs), the SIR averages from -16 dB at 108 MHz to -24 dB at 684 MHz. For a -4 dB downstream PSD attenuation, the SIR averages from -12 dB at 108 MHz to -20 dB at 684 MHz. So decreasing the PSD level in the FDX downstream directly increases the SIR dB for dB.

Figure 29 shows an FDX network model we simulated of the Downstream CM Signal/Interference Ratio vs FDX Downstream Band PSD Boost. The top pair of graphs are for a +4 dB boost; the middle pair of graphs are for no boost or a linear uptilt PSD; and the bottom pair of graphs are for a -4 dB attenuation. The left side shows all upstream and downstream signal levels, downstream echoes, and leakage across the FDX directional coupler ports, and the total combined interference level into the downstream received signal at the CM. The right side shows the resulting



**Figure 28 – Upstream Node Signal/Interference Ratio vs FDX Downstream Band PSD Boost (Top: +4 dB Boost, Middle: No Boost – Linear, Bottom: -4 dB Attenuation)**





**Figure 29 – Downstream CM Signal/Interference Ratio vs FDX Downstream Band PSD Boost (Top: +4 dB Boost, Middle: No Boost – Linear, Bottom: -4 dB Attenuation)**

calculated downstream Signal-to-Interference Ratio (SIR) of the CM. These results show the direct correlation of the downstream node PSD level in the FDX band and the resulting downstream SIR of the CM in that band.

Unlike the node where the full duplex upstream and downstream signals are present in the same spectrum in the FDX band, with the high downstream launch power in the same spectrum as the upstream received signals render the SNR in negative territory, the CM operates in FDD mode. In this case, one of the three FDX sub-bands receives the downstream signal while the other two sub-bands are transmitting upstream, causing adjacent noise and emissions. And, there are echoes into the downstream receiving sub-band resulting in a positive (but low) SNR. Thus the echo cancellation in the node must have a cancellation depth of the interference in excess of 50 dB, while the CM, without self-induced co-channel interference instead requires a cancellation depth of adjacent channel interference of only around 20 dB. The details of such an analysis can be found in [1].

For a +4 dB downstream PSD boost, the SIR averages from 23 dB to 31 dB across the FDX band. For a 0 dB downstream PSD boost (strictly linear without step-downs), the SIR averages from 20 dB to 27 dB across the FDX band. For a -4 dB downstream PSD attenuation, the SIR averages from 16 dB to 24 dB across the FDX band. So decreasing the PSD level in the FDX downstream directly decreases the SIR approximately dB for dB for the CM.

This end-to-end FDX system analysis demonstrates the trade-off one can make for improving the reception on the CM side while degrading the reception on the other (node) side, and vice versa, using a different stepped PSD downstream profile at the node, tailored to the desired directional bit-loading increase while lowering the PSD slightly in the highest frequencies to maintain a constant total composite power in the node. The choice depends on the objective desired – higher spectral efficiency and bit-loading of the chosen signal direction (upstream or downstream) at the expense of the other virtually dB for dB. Note that 3 dB change in SNR corresponds to one bit per symbol difference in bit-loading.

## 9. Conclusion

The transition from conventional sub-split Node + N systems to mid-split Node + 0 systems enables several steps toward the ultimate goal of FDX DOCSIS 4.0 that will significantly increase the upstream bandwidth to multi-gigabit speeds with high throughput. The first step in this transition has been described, starting from conventional sub-split cable systems to mid-split systems to high-split systems and eventually to a phased approach to enabling FDX DOCSIS. The RF spectrum to support both legacy video and data services and the separation of the CMTS core functions that connect to Remote PHY (R-PHY) nodes using a Distributed Access Architecture (DAA) enables this transition in both conventional and Node + 0 networks.

The path to higher speed networks, while supporting legacy video and data consumer premises equipment was explained. Lab and field data was analyzed together to provide the confidence in layering these technologies in an evolutionary fashion. Some mitigation strategies involving signal design launched from the R-PHY node and their effectiveness were explained. Cable Modem architectures to optimize both FDX and legacy DOCSIS signals were explained. Accommodation of legacy video set-top boxes and in-home networking with MoCA in order to fully utilize the 1.2 GHz design bandwidth of our cable systems has been shown.

All of these intermediate steps prepare for the initial introduction of FDX technology in a phased deployment as capacity and speed demands continue to grow. Configurable FDX allocation of RF spectrum will continue to support legacy SC-QAM video, SC-QAM based DOCSIS 3.0, and

OFDM/OFDMA based DOCSIS 3.1 in a compatible and compliant way with an evolutionary path to transition to all IP video with FDX DOCSIS 4.0. Some optimizations of RF spectrum allocation and the conditioning of transmission and reception of FDX signals to improve the effectiveness of the echo cancellation technologies have been proposed. The evolutionary path to 10G networks using FDX DOCSIS 4.0 will enable our networks to grow with the ever increasing demand to support both existing and future services for our customers.

# Abbreviations

ACI	adjacent channel interference
ALI	Adjacent leakage interference
BAU	Business as usual
CACIR	Carrier/ACI Ratio
CCI	Co-channel interference
CIR	Carrier-to-Interference Ratio
CM	Cable modem
CMTS	cable modem termination system
CNR	Carrier to noise ratio
CPE	Customer premise equipment
DAA	Distributed Access Architecture
DOCSIS	Data Over Cable Service Interface Specification
EC	Echo Cancellor
FDD	Frequency Division Duplex
FDX	Full Duplex DOCSIS
FDX-L	A legacy DOCSIS 3.1 cable modem which can transmit in the 108 MHz to 204 MHz Full Duplex upstream channels and receive in the 258 MHz to 684 MHz Full Duplex downstream channels in a high-split access network, or it can receive in the 108 to 684 MHz Full Duplex downstream channels in a mid-split access network with no access to upstream Full Duplex Channels.
HFC	Hybrid Fiber-Coax
IG	Interference Group
MER	Modulation error ratio
MHz	Megahertz
MoCA	Multimedia over Coax Alliance
OFDM	Orthogonal Frequency Division Multiplexing
PA	Power amplifier
PSD	Power spectral density
PHY	Physical
RB	Resource Block
RBA	Resource Block Assignment
R-PHY	Remote Physical (layer)
SC-QAM	Single Carrier Quadrature Amplitude Modulation
SNR	Signal to Noise Ratio
STB	Set-top Box
TCS	Transmit Channel Set
TG	Transmission Group

## Bibliography & References

- [1] Richard S. Prodan, *Full Duplex DOCSIS PHY Layer Design and Analysis for the Fiber Deep Architecture*, 2017 SCTE-ISBE Cable-Tec Expo Fall Technical Forum, Denver, CO.
- [2] SCTE 235, *Operational Practice for the Coexistence of DOCSIS 3.1 Signals and MoCA Signals in the Home Environment*, 2017.
- [3] Multimedia over Coax Alliance, *MoCA 2.0 Field Trial Report*, 2016.

# Bringing Enterprise IoT to Cable

A Technical Paper prepared for SCTE•ISBE by

**John Jason Brzozowski**  
CTO, Vice President Engineering  
MachineQ, a Comcast company  
1800 Arch Street, Philadelphia, PA 19103  
+1-484-962-0060  
[john\\_brzozowski@comcast.com](mailto:john_brzozowski@comcast.com)

# Table of Contents

Title	Page Number
1. Introduction.....	3
2. Definition of Enterprise IoT.....	3
3. Wireless Communication Considerations.....	4
3.1. 2.4 GHz.....	4
3.2. Cellular.....	4
3.3. ISM (Industrial, Scientific and Medical) Bands.....	4
4. Enterprise IoT Platform .....	5
4.1. Integration.....	5
4.2. Security.....	6
4.3. Management .....	6
4.4. Performance.....	6
4.5. Scale.....	7
5. Conclusion.....	7
Abbreviations.....	9
Bibliography & References .....	10

# 1. Introduction

Cable infrastructure is at the heart of voice, video, and Internet services. Subscriberhip across classic cable services has been evolving rapidly across our ecosystem, which prompted MSOs to invest in the development of new products and services that help enable diversification and initiate new subscriber growth.

As such there is burgeoning demand for enterprise class, wireless-based IoT platforms to enable operators to expand their product and service offerings to address the needs of the enterprise, SMB, B2B2C, and even the consumer market segments. Enterprise IoT platforms in particular are essential to enable adoption, when new, disruptive technologies are involved. The availability and use of DIY (Do-It-Yourself) or BYO (Bring-Your-Own) to support the large-scale enterprise IoT deployments is proving to be challenging, if not an impediment. Forcing or expecting adopters to locate, validate, deploy, and manage all the critical components, securely, slows adoption and inevitably growth. Further, Enterprise IoT platforms are particularly important when adopters require the ability to manage IoT offerings across multiple media, wireless or otherwise. As the IoT landscape evolves beyond devices that leverage broadband connections, facilitating the deployment of and offering a seamless, streamlined customer experience is difficult if not impossible if an operator is relying on a swivel chair approach. Enterprise class IoT platforms must offer the ability to manage the deployment and operations of a vast, virtual IoT network that span multiple geographic regions. A purpose built enterprise IoT platform must truly be an enabler of mission critical use cases across a wide range of verticals.

The vision outlined here is one that envelops the following:

- An Enterprise IoT Platform that enables global adoption and enablement
- Spans multiple media, mainly wireless
- Illustrates the need for scale and performance
- Couples in enterprise-class functionality
- Leverages best in class security

The vision components listed above are the core characteristics of any enterprise IoT platform that intends to fuel adoption and operation at scale, while enabling a broad base of market verticals -- from Quick Service Restaurant (QSR), to asset tracking pharmaceutical manufacturing and laboratories, all while spanning the use of multiple wireless media including LoRAWAN, BLE/Bluetooth, and others. Supporting multiple connectivity methods helps to uniquely address each vertical's specific challenges and operating environments.

The world is changing for MSOs, and so is the definition of IoT. More specifically, Enterprise IoT. How we plan for and build to support, manage, and operate Enterprise IoT is rapidly evolving, in lockstep with the demands set forth by our customers. These requirements range from touchless, automated solutions to enable temperature/humidity monitoring and high value asset tracking, to energy monitoring and management for cable infrastructure.

## 2. Definition of Enterprise IoT

The IoT (Internet of Things) carries a wide range of meanings and the term continues to evolve on practically a daily basis. For the purposes of this document, IoT will generally refer to the act of connecting or enabling the connection of a device to utilize the Internet or network, with the intent of interacting with a human or consumer. Or, according to Oxford Languages:



“The interconnection via the Internet of computing devices embedded in everyday objects, enabling them to send and receive data.” [1]

Conversely, Enterprise IoT refers to the enabling of machines to interact and exchange data with other machines in a manner that may or may not be usable by a human. Further, the machines in the case of Enterprise IoT are not strictly required to have a direct connection to the Internet. However, the machines may leverage other elements to facilitate communications.

The Enterprise IoT is viewed by some [2] as the “next advancement in technology that enables physical ‘things’ with embedded computing devices (tiny computers) to participate in business processes for reducing manual work and increasing overall business efficiency.”

An example of consumer-grade IoT, generally, is a video camera connected to the Internet or home network using Wi-Fi via a broadband connection that supports two-way video and audio. Whereas, Enterprise IoT is best exemplified by a wireless (not necessarily Wi-Fi) temperature and humidity sensor that transmits data periodically from multiple locations in a QSR to automate food monitoring. The bandwidth, connectivity model, wireless characteristics, and power consumption vary widely across these and numerous other scenarios.

### **3. Wireless Communication Considerations**

Not all wireless communication protocols are created equal, nor do they need to be. Further, it should not be expected that a single wireless communication protocol be able to universally address every technical and/or business use case. Each form of wireless communications carries its own distinct pros and cons. A summary is provided here, as an in-depth compare-and-contrast of IoT wireless protocols is beyond the scope of this document.

#### **3.1. 2.4 GHz**

Wi-Fi, BLE/Bluetooth, Zigbee all commonly support wireless communications using the 2.4 GHz frequency band. Characteristics include:

- Capable of high bandwidth, two-way communications
- Limited range, limited mobility
- Reasonable propagation for short range wireless communications
- Higher power consumption

#### **3.2. Cellular**

- Capable of high(er) bandwidth, two-way communications
- Extend range and mobility
- Reasonable propagation
- High power consumption

#### **3.3. ISM (Industrial, Scientific and Medical) Bands**

- Capable of low bandwidth, limited two-way communications
- Extended range, mobility support varies
- Optimal propagation for long range wireless communications
- Low power consumption

Given the nature of this document the above focuses on the technical considerations for each type of wireless communications protocol. Cost, while of out scope, is most certainly an essential detail that must be factored into any wireless protocol evaluation.

## 4. Enterprise IoT Platform

Most every wireless communications platform consists of one or more of the following:

1. Wireless devices including sensors or endpoints that typically send and/or receive data
2. Wireless base stations or gateways that enable or facilitate communications to and from wireless devices or remote systems
3. Wireless controllers and related sub-systems that govern the wireless communications between wireless devices and base stations

Common misconceptions suggest that one or all of the above adequately provide the tools required to enable Enterprise IoT for an operator or adopter. In fact, regardless of the wireless communication chosen, the attributes listed above only provide the building blocks to enable basic wireless communication for an Enterprise IoT environment. A platform approach supporting any form of wireless communication must minimally account the following:

- Heterogeneity regarding wireless communications as well as from a device and base station perspective
- Operationalization including measurements and monitoring ensuring service quality and reliability

The following sections further delve into key factors that must be considered when evaluating and/or building a platform oriented for wireless, enterprise IoT.

### 4.1. Integration

Deep integration is essential to ensure the effectiveness of any platform that will effectively be utilized to manage large scale, complex deployments. While integration with internal operational support systems (OSSs) and business support systems (BSSs) is critical, the context here refers to integration with the core hardware elements that truly enable large scale adoption. Wireless communication integration can occur in multiple ways.

First and most common is adherence to a common wireless communication standard. In the case of LoRAWAN® [3], this would imply that a given device or sensor, gateway or base station, and LoRAWAN Network Server (LNS) have all implemented compatible versions of the communication specifications. Compliance at this level simply ensures that elements that are attempting and expecting to interact, can do so.

Alternatively, deeper integration across this confluence of key elements -- where software meets the embedded systems (or hardware) -- affords adopters a “surgical” level control and flexibility. Both matter to confidently, reliability, and seamlessly managing enterprise IoT deployments. This inevitably requires that some form of a platform’s software components reside natively on critical embedded systems. Embedded system integration, in this case, specifically refers to the integration of platform functionality that spans sensors as well as base stations.

## 4.2. Security

Credible enterprise IoT platforms must tout the ability to secure communications end-to-end. While many of the individual elements available today that can be used as piece parts to orchestrate an enterprise IoT deployment are secure, component-level security does not assure that the concatenation of those elements carries the same level of end-to-end security. Nor does this guarantee that critical aspects have not been overlooked. End-to-end visibility across an entire platform provides security assurance specific to data integrity and communications. Data integrity ensures that the data received is, in fact, the data that was sent, while data communications ensures that the transport mechanisms carrying the enterprise IoT data in uncompromised. Further, control of data encryption by the originator (or customer) from a sensor that is independent of transmission or delivery (by an operator) extends a level of independence that secure, trusted enterprise IoT platforms are expected to support. The capability to securely provision sensor communications while maintaining trust must not be overlooked.

In order to maintain the integrity and value that a secure wireless communications protocol affords it must be coupled with a platform that directly enables its streamlined use. Manual configuration and “swivel chair operations” create process gaps that inevitably result in security challenges.

## 4.3. Management

Enterprise IoT management, broadly speaking, builds strongly but not solely on the concept of integration and is the backbone for security. The depth of integration is directly proportional to an organization’s ability to manage key aspects of their Enterprise IoT deployment. Key considerations include:

- Remote, over-the-air firmware management for both sensors and base stations
- Security management and monitoring
- Configuration management for wireless IoT communications
- Base station backhaul connectivity and configuration management

In a world where new vulnerabilities and viruses are discovered on practically a daily basis, the ability to remotely update your IoT infrastructure on a scheduled, automatic basis is essential to ensuring that your deployment is robust. Updating firmware and configurations over the air (or over any backhaul network) can be a complicated affair. Further, with the rapid evolution across the wireless communication spectrum, where the likes of 5G, Citizens Band Radio Service (CBRS), and Narrowband (NB) IoT are becoming a reality, it is perhaps a forgone conclusion that adopters will require a unified platform to enable the management across any and all wireless communication protocols.

Finally, where adopters are choosing LPWAN alternatives like LoRAWAN, given its favorable performance characteristics over long distances, superior propagation, and optimal power consumption, there are in fact bandwidth considerations that must be accounted for from a management perspective. Namely, facilitating firmware updates over the air for LPWAN must be implemented diligently. In an Enterprise IoT sense, diligence means leveraging partial or delta updates in lieu of full image updates, as a means to dramatically reduce update times, positively impact power consumption, and improve hardware life expectancy.

## 4.4. Performance

An enterprise IoT offering that cannot perform well impedes adoption and growth. The advent of inexpensive, battery powered devices that can last for 5 to 10 years (or more) introduces many new performance challenges. The need to provision large volumes of inexpensive devices with a lengthy

lifetime translates into millions if not billions of sensors that need to be managed over time. This includes high performance provisioning to support rapid deployments and the processing of data from the same. Processing data from an exponentially growing population of enterprise IoT sensors has the potential to yield billions of data payloads -- daily.

Succinctly put, solving problems, creating value, and generating revenue are tightly coupled with the rate at which sensors and base stations can be provisioned and deployed. One of the main purposes of any high-performance enterprise IoT platform is to enable the currency of enterprise IoT, data, to flow as quickly and in as timely a manner possible. The frequency of data transmissions is often a byproduct of the market vertical or use case.

Essential characteristics related to Enterprise IoT performance include:

- Web-scale Application Program Interfaces (APIs) and backend implementations that support large volumes of provisioning events that enable sensors instantaneously
- Intuitive user interfaces allowing for streamlined, high-performance management
- Carrier-grade base station performance and sensor data processing ensuring that payload delivery performance is optimized globally

## **4.5. Scale**

Global scale and locally optimized performance and resiliency are absolute musts for modern day enterprise IoT adopters. Enterprise IoT adopters today have different characteristics and requirements than consumer IoT adopters. Enterprise IoT adopters are multi-national, global corporations that are seeking to form partnerships and build relationships with similarly-sized firms that are able to service their needs around the globe, not just in one region specifically. The scope of scale reaches far beyond, but is anchored in, the scale and performance of an enterprise IoT platform. Scale encompasses service (installation, deployment, etc.), support, and global capacity.

Global scale is, in many ways, a culmination of the characteristics outlined throughout this document. Assembling a platform and business to address the evolving needs of enterprise IoT adopters globally starts with borrowing from the experiences in scaling other businesses, including Internet, voice, and video. Scaling Enterprise IoT, while sharing many similarities with scaling some of the largest data, voice, and video networks in the world, carries many unique nuances. Most Enterprise IoT platforms and networks are designed to support enormous volumes of small payloads originating from sensors across a large heterogeneous (global) customer base. Further, ensuring that base stations, which route the lifeblood of most LPWAN networks, have robust, reliable backhaul is essential to every enterprise IoT deployment.

## **5. Conclusion**

Orchestrating a large-scale Enterprise IoT deployment is a complex endeavor. Choosing which wireless communication protocol(s) best address requirements alone is a complex task. Identifying or building a platform that enables an organization to manage it globally increases the sophistication, timing, and investment. Factoring in carrier-class scale, performance, and best-in-class embedded systems like sensors and base stations is also critical to any well-intentioned plan. Ensuring that all embedded systems are qualified regionally, and that all regulatory guidelines are adhered to, are significant tasks.

It is important to not discount the effort involved in operationalizing the use of Enterprise IoT to generate revenue or accelerate cost savings across your own or your customer's enterprise. Implementing an enterprise IoT solution for temperature and humidity monitoring across thousands of locations, globally, or leveraging it to deploy advanced power management and analytics across cable networks for an MSO,

are both real and valid use cases. Discounting the need for and the investment required to build a robust, globally scalable multi-wireless media enterprise IoT platform will impose substantial limitations. Limitations that will impact an operators ability to deliver a reliable, future proofed service that will fail to comprehensively and proactively addresses their customers' requirements. Customers' whose revenue relies on you and yours on them.

## Abbreviations

MSO	Multiple Services Operator
SMB	Small/Medium Business
B2B	Business to Business
B2B2C	Business to Business to Consumer
IoT	Internet of Things
BYO	Bring Your Own
DIY	Do It Yourself
QSR	Quick Service Restaurant
LNS	LoRAWAN ® Network Server
OSS	Operations Support Systems
BSS	Business support systems
5G	5th generation mobile network
CBRS	Citizens Broadband Radio Service
NB IoT	Narrowband Internet of Things

# Bibliography & References

[1] Oxford Languages

[2] <http://www.enterox.com/IoT/articles/enterprise-internet-of-things.htm>

[3] <https://lora-alliance.org/about-lorawan>

LoRaWAN® Specification v1.0.2; LoRa Alliance®

LoRaWAN® Specification v1.0.3; LoRa Alliance®

LoRaWAN® Specification v1.1; LoRa Alliance®

*What is the LoRaWAN® Specification?*; LoRa Alliance®

# **Private Mobile Networks - A New Service Option for Enterprise Wireless Connectivity**

## **Understanding Enterprise Needs and Use Cases for Private LTE/5G Services**

A Technical Paper prepared for SCTE•ISBE by

**Dr. Mehmet Yavuz**  
CTO and co-founder  
Celona, Inc.  
10061 Bubb Rd Ste 300, Cupertino, CA 95014  
(408) 906-8198  
mehmet@celona.io



# Table of Contents

<b>Title</b>	<b>Page Number</b>
1. Introduction.....	3
2. What is a Private Mobile Network? .....	3
3. Determining Enterprise Needs and Technical Requirements.....	4
4. Use Cases.....	6
4.1. Supply Chain and Logistics.....	6
4.2. Healthcare .....	6
4.3. Manufacturing.....	7
4.4. Higher Education.....	7
5. Conclusion.....	8
Abbreviations .....	8

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - CBRS LTE Network Design Tool.....	5

## 1. Introduction

Just in the last few years, a multitude of new technologies and methodologies have emerged that are forcing cable telecom engineers to reevaluate their networking architectures and service offerings, particularly those targeting the enterprise. Whether it is 5G, the Internet of Things (IoT), low-power wireless protocols, open source hardware, or some combination of these new tools, there is little doubt that big changes are coming to wireless service offerings for Multi-Service Operators (MSOs) and cable operators of all kinds. The viable application of these technologies vary widely, and it is up to MSOs to sort near-term practicality from marketing vaporware.

One recent development that has immediate potential for MSOs is the decision earlier this year by the FCC to make the Citizens Broadband Radio Service (CBRS) spectrum band available for use by businesses. The net result is that MSOs/cable operators now have a new option for predictable mobility that can serve a variety of needs on mission-critical infrastructure for enterprises across a broad swath of industries, from outdoor campuses, to healthcare, manufacturing, logistics/transportation, higher education, and more. This new option is called a private mobile network and it is poised to change how MSOs view wireless networks in their environments.

## 2. What is a Private Mobile Network?

First, let's provide some background on what we mean by a private mobile network in the context of enterprises. We all know that enterprises heavily rely on wireless connectivity both indoors and outdoors, two different environments that present separate challenges for different types of wireless technologies. For the last two decades, "wireless" for the enterprise setting meant two things; Wi-Fi or a public cellular network.

The advent of private mobile networks is something new. They are very similar to the public cellular network that most of us use every day with our smartphones and tablets. They are built on technologies like LTE and 5G as the wireless protocols for connectivity, but the networks themselves are not owned and operated by wireless service providers like Verizon, AT&T, and others. A private mobile network is an LTE or 5G network that is owned and operated by a single organization and geographically bound by that company's property (like a smart factory, hospital, university campus, or shipyard). In the U.S. these networks use the CBRS spectrum band between 3.55-3.7Ghz and can be used by enterprises to give them their very own LTE or 5G network, but managed as easily as the Wi-Fi network.

So why add a new wireless option? First, a private mobile network based on CBRS spectrum operates within a dedicated and interference-free spectrum by FCC regulations in the United States. Not only does this significantly reduce the ongoing operational expenses and support costs, it makes a private mobile network ideal for applications that need to be up and running at all times. Solutions like IoT devices and systems need ultra-reliable connectivity in order to perform their core functions and, from a wireless perspective, private mobile networks are the only option capable of delivering on that requirement.

In addition to reliability, CBRS-based networks provide the type of latency that businesses need for both current systems and future applications. Even if based on LTE, a private mobile network can provide specific performance guarantees at an established Quality of Service (QoS). Today's networks deliver 20ms latency one-way. That latency will be even further reduced for private mobile networks based on 5G. This is a critical requirement for any type of factory automation task and for many Industrial IoT applications.

Another benefit of this new spectrum-based “traffic lane” is that it ensures the enterprise data traffic is kept local and separate from networks that may be used by guests or other personnel that do not need access to secure data. That built-in security can be a critical element for business and safety reasons, as well as a big reason that enterprises will avoid services leveraging public cellular networks delivered by the traditional wireless operators.

CBRS-based private mobile networks are also ideal for connectivity across outdoor settings such as university campuses, transportation yards, outdoor warehouse locations, or manufacturing site parking lots. With the ability to cover one million square feet of outdoor space with a single outdoor wireless access point (AP), a private mobile network significantly reduces the amount of outdoor cabling required to support a wireless network infrastructure - reducing a good chunk of an enterprise’s capital expense and thus making a private mobile network service attractive.

Private mobile networks are an ideal extension of another new networking architecture that is quickly establishing itself in the enterprise sector - edge computing. As demands for low latency and computing intelligence increase, most IT network architects are looking for the means to place at least some of this functionality physically closer to applications themselves, thus making the “edge” of the network closer to the “action” so to speak. Analytics, information applications, business processes and other functions will likely still reside in the cloud or in an onsite data center, but private mobile networks can extend operational technologies and thus the network edge to the enterprise for security applications, computer vision apps, system automation, delivering predictable communications for guided vehicles, and guaranteeing reliable voice communications for staff.

### **3. Determining Enterprise Needs and Technical Requirements**

In evaluating the enterprise market for private mobile network services, it’s clear that there are businesses that do NOT yet need a private mobile network. If they are not pushing the limits of their current wireless network or if they are not looking at business process automation or the Industrial IoT, it’s possible they do not have a need for this technology yet.

However, based on the demand we’re seeing here at Celona, there are definitely many horizontal service applications for private mobile networks that are being currently evaluated, so let’s discuss what MSOs need to consider when implementing services based on this technology. There will likely be many different ways of providing private mobile networks, each depending on the specific needs of the organization’s IT department. Some groups will want to have this supplied to them and managed by a third party, like an MSO, wireless operator, or channel partner. Others may wish to simply purchase an end-to-end system and install and manage it themselves.


Regardless of the go-to-market model, in order for private mobile networks to truly fulfill their potential for the enterprise, they need to evolve their basic architectures to leverage the devices and apps that are primarily used “at work” and integrate with existing device and network-level access policies and provide enterprise IT full ownership of data, analytics and insights. Such an architecture would be tightly integrated to the existing enterprise network underlay and deployed as fast as you would deploy an enterprise Wi-Fi solution today with the primary goal of serving apps that demand guaranteed service levels. With a flexible IT-friendly approach, managed private mobile networks can enable custom implementations across islands within the enterprise: on the road, across remote sites, at the branch locations and within the campus - indoors and outdoors. By making such a solution IT-friendly, MSOs can ensure their offering will stand out compared to those from a telecom-oriented background (such as a wireless operator).

One of the biggest issues is figuring out the total number of indoor and outdoor CBRS LTE access points (AP), an enterprise might need at a facility given size, geo-location, density and performance requirements. Let's look at one scenario for an indoor factory and what would be required for establishing total network coverage, but also then designing that network to provide maximum capacity as well.

With the ability to cover 1M SQFT of outdoor space with a single outdoor wireless AP, a private mobile network significantly reduces the amount of outdoor cabling required to support a wireless network infrastructure - reducing a good chunk of your capital expense.

Additionally, if you are designing for coverage, CBRS spectrum availability does not impact the number of APs required in the network. In standard mode of operation, since all traffic flows in the download and upload direction to an LTE radio is scheduled by the infrastructure, even a single channel of 10MHz across all APs will work. If you are designing for network capacity, as you move from the maximum possible 150MHz of CBRS spectrum availability across the entire network to 80MHz - which happens to be the worst case scenario for any environment in the United States - you are looking at about only 10-15% increase in the total AP count in a given enterprise network.

If we consider an indoor facility size about one million square feet, there is a reasonable chance most facilities this size will have around 2500 devices (give or take a few dozen) activity simultaneously. In this scenario, we would require around 43 APs to ensure highly-reliable coverage of the entire facility. If you're designing for capacity, we've included a chart below that details what kind of infrastructure will be needed, assuming a download capacity of 3 Mbps and an uploaded capacity of 1 Mbps.



**Figure 1 - CBRS LTE Network Design Tool**

As you can see, the design scenarios for enterprises to implement private mobile networks are already well-defined and available. If you find this type of tool useful, Celona's network planner available at

<https://planner.celona.io>. There are others that would give you the ability to compare and contrast multiple solutions from different vendors, as any reasonable purchaser would need.

## **4. Use Cases**

Now that we've established enterprise design requirements and general applications, let's take a closer look at specific vertical market use cases for private mobile networks.

### **4.1. Supply Chain and Logistics**

The need for mobility for this industry is real with supply chain facilities going through significant changes in layout and product placement across aisles on an hourly basis. With its ability to “hold onto” the signal in challenging environments, a CBRS-capable device can support a new set of applications that require real-time response from the network. For instance, leveraging computer vision to automatically and wirelessly identify fleet vehicles as they enter and depart warehouse facilities for the purposes of inventory tracking and fleet management.

From my own personal experience, we are working with a major global supply chain firm that requires connectivity within a large, outdoor railyard. Using the public cellular network was not an option since this organization does not want their sensitive private data traveling across the operator network. Implementing a private mobile network - in this case based on LTE - the firm is planning to blanket the area with reliable wireless connectivity for video applications without breaking the bank.

The rise of private mobile networks are an inevitability for supply chain firms, but there will likely be many different ways of implementing them, depending on the specific needs of the supply chain management IT department. A solid percentage of these organizations will want to have this supplied to them and managed by a third party, including MSOs. The opening of the CBRS spectrum in the United States - with many variations of it on the horizon across different countries - has made private mobile network services for this market segment a real opportunity for many years to come.

### **4.2. Healthcare**

Private mobile networks based on LTE and eventually 5G represent an entirely novel means of healthcare IT professionals delivering the predictable performance that modern medical facilities require in wireless communications. With security, privacy and reliability as core tenants, CBRS based LTE is available today to meet the connectivity needs for iOS and Android devices operated by the clinical staff.

As a result, healthcare institutions, including hospitals, clinics, campuses and even testing centers are actively evaluating private mobile networks as a potential solution for their wireless needs. The current pandemic has also uncovered new use cases that would require the deployment of private mobile networks using LTE or 5G in order to augment their existing wireless connectivity options. Beyond general enterprise connectivity, private LTE/5G opens the door to supporting mission-critical applications deployed via enterprise-owned devices on their own clean and dedicated lane of wireless communication. Thanks to CBRS spectrum availability in the United States via FCC's recent commercialization of the 3.5-3.7GHz spectrum, private mobile networks can enable critical mobile connectivity indoors and outdoors across hospitals, newly constructed make-shift healthcare treatment locations, and drive-through testing centers.

A typical drive through testing center like the one at 945 Sansom St. in Philadelphia covers an area of approximately 60,000 sq feet, which can be covered with a private LTE network of just 3 access

points. Due to its strong rate vs range performance, deployment and management of the radio infrastructure can also be accelerated. Private mobile networks, by design and per standard, come with configuration of service level objectives: latency, throughput, packet error rate and jitter metrics can be pre-defined against a given set of mission critical devices and/or applications – making sure that what’s most important is protected.

### **4.3. Manufacturing**

The manufacturing market segment often has connectivity needs similar to transportation and supply chain organizations, as both share large swaths of “non-carpeted” and outdoor facilities that require connectivity. These include vehicle-mounted devices, as well as other “ruggedized” devices including smartphones and tablets carried by staff. These devices require inference-free operation, so dedicated spectrum for devices from Samsung, Zebra, and Getac can be critical. In a factory setting, this may also include automated guided vehicles carrying out mission-critical tasks for which other solutions may not be reliable enough.

Manufacturers may also evaluate private LTE and private 5G networks for latency-sensitive operations that may include predictive maintenance and factory automation. These applications are often running non-stop twenty-four hours a day while simultaneously requiring real-time connectivity and response, making them ideal for the reliability and low-latency provided by 5G and LTE.

Many manufacturers are also evaluating private mobile networks for applications that include video surveillance for security. These video applications could also extend to computer vision cameras for inventory and other “parking lot” technology that might be utilized outside the manufacturing plant.

Finally, many manufacturers require interference-free connectivity for executive staff mobile devices and dedicated spectrum for new product lines that are currently under research & development that, due to their experimental nature, cannot be part of the existing Wi-Fi network.

### **4.4. Higher Education**

Institutions of higher education have a unique set of needs that set them apart from standard enterprises when it comes to network architectures. Large amounts of outdoor areas, 24/7 operation of computer labs and research facilities, a high reliance on voice communications between staff members, and municipal-like requirements for parking lots, emergency response systems and video surveillance. These use cases combine to create multiple challenges for predictable performance for wireless networking - especially when the student Wi-Fi is the most important application for many and cannot be interrupted at any time.

Up until recently, higher education IT departments had two choices for outdoor wireless - Wi-Fi or public cellular networks. The FCC’s recent decision to make the CBRS spectrum available for use in private mobile networks has now provided a novel solution for outdoors for critical infrastructure connectivity - university’s very own LTE and 5G wireless infrastructure.

Given the coverage patterns of a CBRS LTE access point that reach up to million square feet, IT departments in higher ed can extend critical infrastructure connectivity across the campus without breaking the bank by digging up campus parking lots and other property to lay more fiber. Some of the outdoor use cases that our university customers are evaluating include public safety cameras, emergency phones, Internet backhaul for shuttle buses, outdoor lighting controls and monitoring solutions. Indoor use cases include everything from lighting controls and door locks to vending machines and panic buttons.

A big point for the higher education is that these IT departments want to extend reach to areas where Wi-Fi is not a fit and to keep Wi-Fi spectrum open for users and saving private LTE or 5G for infrastructure.

## 5. Conclusion

The bottom line is that private mobile networks offer a robust solution for multiple near-term enterprise needs. In an era in which we are assailed with far-off applications and nice-to-have functionality, private mobile networks deliver real performance with practical value today. If you're an MSO reading this now, I can also tell you that I am actively engaged with many of your likely competitors – from wireless operators to cloud service providers - as they are carefully evaluating different solutions from the multiple vendors in this market segment. What's clear is that for the enterprise, private LTE and 5G networks are going to be a critical option for enterprise connectivity going forward. If you plan on serving this market segment, you owe it to yourself to learn more.

## Abbreviations

AP	access point
CBRS	Citizens Broadband Radio Service
IoT	Internet of Things
MSO	Multi-Service Operator

# **In Pursuit of the Dark NOC: Driving Change With Automation & AIOps**

A Technical Paper prepared for SCTE•ISBE by

**Marcus Rebelo**

Head of Americas Sales Engineering

Resolve Systems

<https://resolve.io>

775.842.4469

[Marcus.rebelo@resolve.io](mailto:Marcus.rebelo@resolve.io)



# Table of Contents

Title	Page Number
1. Introduction.....	3
2. Key Challenges Facing CSPs.....	3
3. The Role of AI and Automation in the Dark NOC .....	4
4. Building Culture to Embrace Automation and AI .....	5
4.1. New Roles in the New NOC.....	5
4.2. Crafting a Culture of Automation.....	5
4.3. Capturing Tribal Knowledge with Automation .....	6
5. Jumpstarting & Scaling Network Automation .....	6
5.1. Proactive Network Testing .....	6
5.1.1. A Real-World Proactive Network Testing Example .....	7
5.2. Incident Resolution .....	7
5.2.1. Real-World Incident Resolution Examples .....	8
5.3. Network Provisioning .....	8
5.4. Leveraging AIOps.....	8
6. A Blueprint for Getting Started – Actionable Steps to Transform the NOC.....	9
7. Summary – Innovation Powering New Services.....	10
Bibliography & References .....	10

# 1. Introduction

Network operations have evolved radically in the wake of digital transformation – and the increasing infrastructure complexity that accompanied it. As communication service providers (CSPs) roll out and support next-generation technologies like NFV, SD-WAN, IPv6, 5G, and soon 6G, network engineers face a myriad of challenges and changes in their daily operations.

The coronavirus pandemic has presented additional challenges, such as new network traffic patterns and a shifting workforce and has put unprecedented pressure on network teams to safeguard business continuity, network security, and quality of service at all costs. Reliable and high-performing connectivity has truly never been more important, making the role of the NOC more critical than ever before.

This paper explores the future of network operations, examining where we are and what's next. Drawing on real-world experience, the paper presents a blueprint for success when it comes to modernizing network operations. It also delves into the increasing role that automation and AI for IT Operations (AIOps) play in laying the foundation for the next-generation of NetOps, as well as how they are helping CSPs address current IT challenges including those stemming from COVID-19.

Becoming an “operator of the future” will require CSPs to maximize the potential of artificial intelligence, machine learning, and automation in pursuit of the so-called “Dark NOC.” There are a number of actionable steps that teams can take today to achieve both quick wins and long-term success on the road to the Dark NOC and beyond.

## 2. Key Challenges Facing CSPs

CSPs face extraordinary circumstances when it comes to network operations due to the immense scale of the infrastructure they manage. New technologies are aggressively ramped up to remain competitive and meet customer demands, resulting in much larger scale deployments compared to average enterprise organizations – whether it's network virtualization, software-defined infrastructure, IPv6, 5G, or the imminent 6G.

Key challenges facing CSPs include:

- Managing network infrastructure complexity, including legacy networks, while rolling out complicated and demanding new technologies
- Maintaining high levels of network connectivity and service delivery
- Contending with massive amounts of data being generated by the infrastructure and the monitoring tools that have been deployed to keep an eye on it
- Ensuring network security
- Morphing, unpredictable network traffic patterns
- Meeting tight margins, requiring continual improvements in operational efficiency
- Quickly rolling out new services and infrastructure (without increasing headcount) to remain competitive

As we battle a pandemic, CSPs must also contend with workforce fluctuations (or in some cases reductions) and tightening budgets, meaning that NOC technicians need to do more with less and workloads need to be left-shifted.

To address these challenges, CSPs must embrace new technologies and advance their use of automation, AI, and machine learning. Manual processes simply cannot keep pace with the evolution, complexity, and rapid change in modern networks.

### **3. The Role of AI and Automation in the Dark NOC**

Automation and AI are front and center in the Dark NOC. Today, automation handles not only the mundane, repetitive tasks that most network engineers dread, it also tackles complex processes that are otherwise impossible to execute on a regular basis due to time constraints. In the future, automation will replace even more of the manual activities performed by network engineers. Humans will focus on new technologies, creative problem solving, and innovation, while automation addresses much of the day-to-day.

While the Dark NOC is the long-term goal, automation offers immediate, incremental benefits, allowing operations to run much more efficiently with fewer resources. In fact, several CSPs have consolidated their NOCs with just a modest level of automation, resulting in significant costs savings. Others are turning to automation to fill the gap as baby boomers in the NOC reach retirement age.

AIOps is another important technology powering the Dark NOC. It harnesses artificial intelligence and machine learning to ingest, aggregate, correlate, and analyze millions of data points to produce insights into the health and performance of the network infrastructure and applications running on top of it. This powerful technology dramatically reduces alarm noise by performing advanced event correlation to highlight real problems and intelligently group events to isolate the root cause of issues. AIOps also has the capacity to predict and prevent outages by identifying anomalies, conducting advanced pattern analysis, and performing intelligent, dynamic thresholding. Collectively, these capabilities improve performance and reliability, accelerate incident response, and prevent outages, all while streamlining operations.

Leading analyst firms have reported a significant rise in automation and AIOps initiatives in the last year, laying the groundwork for the future state of the NOC. Gartner research indicates that 94% of executives are investing in I&O automation or plan to start. And, while less than 20% of the Global 5000 have a centralized automation function today, that is expected to grow to 90% by 2025.

EMA Research reports that 85% of enterprises have AIOps underway or planned as a major initiative. Their data shows that AIOps in deployment correlates strongly with more progressive levels of automation, closely coupling these technologies. In fact, combining the two offers a closed loop of discovery, analysis, detection, prediction, and automation.

Together automation and AIOps offer a path to the Dark NOC as the insights from artificial intelligence fuel the automation ecosystem, with increasingly less oversight from humans and increasingly more direction from artificial intelligence. These capabilities bring us closer to the promise of self-healing networks and the desired state of the future NOC.

## 4. Building Culture to Embrace Automation and AI

Creating the Dark NOC requires not only great technology, but also a culture that embraces automation and AI, a framework for organizational development, and process optimization.

### 4.1. New Roles in the New NOC

The first and most significant step in advancing the NOC doesn't involve technology at all. It's all about people, organizational structure, and up-leveling current employees into new positions that are focused on leveraging automation and AI to solve problems and improve efficiency.

Some network engineers may worry about the impact of automation on their jobs. The reality is that networking roles are moving in an exciting direction with new opportunities to drive the business. The goal is to up-level and retrain, not to displace. AT&T set an early example in 2016 by investing \$1 billion in workforce retraining specifically to pivot to a software-centric network.

Among decision makers at firms that are adopting automation, 26% told Forrester that they face challenges with culture and change management; 25% believe they have gaps in their organizational structure, alignment, and readiness; and 25% said their firm lacks an overall vision or strategy for automation. Similarly, Gartner asked clients to identify their top three organizational challenges related to automation: 53% reported a shortage of people with necessary skills; 46% cited a lack of documentation of existing processes; and 44% said cultural resistance was a concern.

Taking the time to define what your Dark NOC team looks like and crafting new roles and responsibilities is essential for success. This should be part of a broader effort beyond the NOC to create, as Forrester puts it, “an automation strike team” and Center of Excellence (CoE).

Repurposing skillsets to support the transition to a Dark NOC means shifting people to focus on more data-driven customer service and moving them up in the knowledge stack to identify candidates for automation, validate automations, build automation content, and take on responsibilities that require more human-centric skills. What have traditionally been operations-oriented roles are morphing into more critical-thinking roles to facilitate digital transformation. New skillsets require understanding data and processes and making thoughtful decisions about how to leverage that data to optimize operations, advance services, and deliver innovation.

More specifically, emerging roles in the NOC include network automation architects who focus on determining the best processes to automate and how to go about doing so, as well as network automation managers who establish governance processes, run automation teams, manage the toolsets, and collaborate with business stakeholders.

### 4.2. Crafting a Culture of Automation

With the right roles in place, CSPs can implement initiatives that help promote the culture of automation both within and beyond the NOC. Some real-world examples include:

- Educating the workforce on the benefits of automation and how they can be a part of it
- Identifying champions and giving them a platform to socialize the benefits of automation
- Creating an internal brand for your automation strike team

- Providing easily accessible training on automation and AI tools
- Developing an automation forum to share ideas and promote what's coming down the automation pipeline
- Publicizing how automation is contributing to strategic initiatives and how it aligns with business objectives
- Showcasing automation success and rollouts in internal newsletters, videos, town halls, internal chat channels, and employee portals
- Establishing a simple process for people to recommend processes for automation
- Incentivizing employees to participate with automation awards and cash prizes

### **4.3. Capturing Tribal Knowledge with Automation**

All too often tribal knowledge and expertise still reside primarily with a handful of subject matter experts (SMEs), especially when it comes to complex technologies and legacy toolsets. Documenting this know-how (and best practices in general) is critical to advancing automation initiatives and should be part of any plan for long-term success.

Valuable tribal knowledge can be captured and encoded into automation workflows, preserving this expertise for the long term and making it available in such a way that it scales. SMEs can even capture knowledge with conditional logic to create interactive automations that step through best practice procedures while automation executes the tasks behind those steps. In doing so, CSPs can left-shift workloads, minimize unnecessary escalations, and keep costs down.

## **5. Jumpstarting & Scaling Network Automation**

Forrester reports that many networking teams are still in the early stages of automation, despite advances in virtualization and programmable infrastructure. One of the key challenges to any automation initiative is determining where to start.

The use cases below provide a guide to jumpstarting and scaling network automation initiatives to support digital transformation and next generation technology rollouts, as well as to improve operational efficiency. Rolling out these use cases at scale will put CSPs well on their way to achieving the Dark NOC.

### **5.1. Proactive Network Testing**

Automating proactive network tests can avoid costly and time-consuming truck rolls, prevent network outages, and provide real-time insights into service quality – not to mention reducing repetitive workloads on the NOC.

Automation can be leveraged to:

- Execute multi-step testing processes, like PIM testing, circuit testing, and customer turn-ups
- Capture network testing and diagnostic information for analytics, compliance, auditing, and change management purposes

- Build comprehensive reports on the health of the network, so key stakeholders and executives can always be on top of service delivery and quality

Benefits of automating network testing include:

- Identifying performance and connectivity issues before they impact users
- Speeding up network testing, while simplifying the process
- Reducing reliance on operator interpretation
- Improving efficiency by replacing manual efforts with automation
- Enforcing consistency and standardization across testing procedures

### **5.1.1. A Real-World Proactive Network Testing Example**

Cell tower outages were causing major impacts for the customers of a large CSP, and their NOC needed a way to validate customer service complaints, perform diagnostics, and quickly implement fixes. With a daily rate of 500K alarms coming in, their Tier 1 analysts couldn't keep pace.

They also needed to automate a set of complex testing procedures on more than 57,000 cell towers. By running these procedures consistently, they could verify mobile connectivity for their millions of subscribers.

Today, an automated process polls the entire radio network every hour and compares outages it identifies to existing tickets and events. It then automatically updates these tickets and generates validated events that are sent to the NOC for follow-up. Additionally, complete coverage and outage reports are auto-generated and sent to each of the market owners, providing unprecedented and accurate visibility into real-time network health and performance.

By automating network testing procedures, the CSP saves 40,000 man hours every year while also accelerating incident resolution and improving service delivery.

## **5.2. Incident Resolution**

Network uptime and performance are mission critical, putting incredible pressure on the NOC to resolve incidents as quickly and efficiently as possible. Automation can radically transform the entire incident resolution process and overcome common challenges – for example, avoiding unnecessary escalations or delays by encoding systems access in the automation itself, thereby enabling lower level analysts to resolve the problem on their own.

Automation can be leveraged to:

- Reduce alarm noise by automatically validating events, correlating related events, and consolidating duplicates
- Collect troubleshooting data from multiple systems across environments and locations, build timelines, and diagnose issues
- Centralize incident management to oversee all aspects of the incident resolution process from a single pane of glass
- Automatically resolve common incidents with pre-built workflows that can be triggered in a variety of ways

- Arm admins with interactive automations to address outages that require more complex resolution workflows, including step-by-step instructions, incremental automations, and decision paths
- Enable agents to safely execute remediation tasks without system access or the necessary CLI or coding skills
- Eliminate time-consuming handoffs between the ITOps and NetOps teams with centralized orchestration of the end-to-end, service-level workflows
- Auto-create and update tickets with a log of all actions and automations that were completed

### **5.2.1. Real-World Incident Resolution Examples**

The NOC for a leading multinational CSP leverages automation to process both traditional and non-traditional alarms, such as network configuration, compliance, and inventory validations. By automating alarm handling, trouble ticketing, triage, intelligent decision-making, and dispatch operations, the organization has accelerated incident resolution time from more than 1800 minutes to less than 60 seconds.

Similarly, another North American provider also relies on automation to reduce alarm noise, enrich events, and triage resolution procedures. They have even extended automation to customer communications. When nodes are impacted by an outage, an automated process quickly cross references those nodes with their CRM system to determine which customers are impacted, and notifications are automatically sent out to keep customers informed in real time.

## **5.3. Network Provisioning**

Automating network provisioning is significantly faster, more reliable, and more accurate than manual efforts – and automation eliminates human error. It also enables CSPs to deliver scalable new services to customers much quicker.

Automation can fast-track network provisioning by:

- Executing pre-provisioning checks to eliminate hours of manual effort
- Conducting pre-configuration checks for authorization and/or billing
- Automating post-configuration checks and any necessary remediation steps to verify connectivity and performance
- Performing network configuration tasks to enforce standard operations and secure configurations
- Auditing every process – whether human-directed or automated – for a real-time, centralized audit trail essential for troubleshooting and compliance
- Integrating and automating updates with service and change request systems

## **5.4. Leveraging AIOps**

AIOps and automation can radically improve network reliability and reduce outages by combining predictive analytics with proactive automations.

Key applications include:

- **Sequential Pattern Analysis & Time-Series Event Correlation:** AIOps has the ability to normalize and sequence millions of events across applications and infrastructure into a time series that can be

analyzed by machine learning algorithms to identify patterns. These patterns can then be leveraged to identify potential outages and trigger intelligent automations to resolve problems proactively.

- **Dynamic Thresholding & Multi-Variate Anomaly Detection:** Anomaly detection algorithms use unsupervised machine learning to learn network environments over time, recognize expected behavior, and set dynamic thresholds across multiple performance metrics that account for seasonality. As events are analyzed in real time, they are compared to expected behavior to provide alerts and/or trigger automated actions when a sequence of events demonstrates anomalous activity.
- **Dynamic Capacity Adjustments:** AIOps leverages historical utilization trends to predict when infrastructure will become non-operational due to capacity exhaustion. These predictive capabilities ensure more capacity can be added dynamically via triggered automations (or through manual intervention) to avoid outages stemming from common capacity issues.

## 6. A Blueprint for Getting Started – Actionable Steps to Transform the NOC

The technology to implement a Dark NOC is rapidly maturing. Depending on where you are in the process, here are some recommended steps towards achieving self-healing networks and automation excellence.

Phase One (2-3 months):

- Document key processes and troubleshooting activities
- Improve visibility with discovery, network mapping, and identifying application flows across hybrid networks
- Identify quick wins for out-of-the-box automation use cases
- Validate integrations across the technology stack
- Identify internal champions for automation and AIOps initiatives
- Build executive support

Phase Two (3-6 months):

- Identify the next phase of automations, including more complex processes
- Optimize the processes on paper before automating them
- Leverage technologies that deliver noise reduction, event clustering, and correlation
- Engage employees in identifying processes to automate

Phase Three (6-12 months):

- Create an Automation Center of Excellence to centralize automation efforts
- Connect AIOps to automation and start with an initial set of autonomous actions
- Activate your “army of automators” to scale
- Integrate analytics and automation
- Expand usage of autonomous automations that can be triggered by AIOps insights



## 7. Summary – Innovation Powering New Services

Evolving NetOps means addressing operational challenges at scale. It requires a deep understanding of how the NOC will evolve, as well as how people, processes, and technology intersect. It requires leveraging automation and AIOps to manage increasing network complexity that has far surpassed the capacity of manual processes and human analysis.

By darkening the NOC, CSPs can address tactical challenges related to cost reduction, workforce fluctuations, and tight margins, but even more importantly, they can free valuable resources to focus on innovation and creative uses of technology to offer new and improved services to propel the business forward.

## Bibliography & References

“Building Our Software-Centric Workforce,” *John Donovan*; AT&T, [https://about.att.com/innovationblog/building\\_workforce](https://about.att.com/innovationblog/building_workforce) (accessed August 13, 2020)

“COVID-19 Accelerates Enterprise Use of Automation in Digital Transformation,” *Brian Solis*; CIO, <https://www.cio.com/article/3562697/covid-19-accelerates-enterprise-use-of-automation-in-digital-transformation.html> (accessed August 13, 2020)

Data-Driven IT Automation: A Vision for the Modern CIO, *Dennis Nils Drogseth, Dan Twing*; Enterprise Management Associates, April 2020

Jump-Start Your Network Automation, *Andre Kindness, Chris Gardner*; Forrester Research, 6 February 2020

Market Guide for Service Orchestration and Automation Platforms, *Manjunath Bhat, Daniel Betts, Hassan Ennaciri, Chris Saunderson*; Gartner, 17 April 2020

“The Role of AI in Creating Energy Efficient Autonomous Networks,” *Frank Rayal*; Mobile World Live, <https://www.mobileworldlive.com/huawei-updates/the-role-of-ai-in-creating-energy-efficient-autonomous-networks/> (accessed August 13, 2020)

Seven Shifts That Will Define The Next-Generation Networking Professional, *Andre Kindness, Christopher Voce, Renee Taylor, Shayna Neuburg*; Forrester Research, 11 April 2019

# Is “Unity Gain” Still the #1 Objective?

**Maybe YES!**

A Technical Paper prepared for SCTE•ISBE by

**John Ulm**

Engineering Fellow, Broadband Systems  
CommScope – CTO Network Solutions team  
Moultonborough, NH 03254  
+1 (978) 609-6028  
john.ulm@commscope.com

**Dr. Zoran Maricevic**

Engineering Fellow  
CommScope – CTO Network Solutions team  
15 Sterling Drive Wallingford, CT 06117  
+1 (203) 303-6547  
zoran.maricevic@commscope.com

**Dr. Frank O’Keeffe**

Distinguished Systems Engineer  
CommScope – CTO Network Solutions team  
2 Johnstown Close, Glounthaune, Cork, Ireland T45V094  
(+353) 87 668 8789  
frank.okeeffe@commscope.com

# Table of Contents

Title	Page Number
1. Introduction.....	4
2. Unity Gain Overview .....	4
3. Network Capacity Planning for Cable 10G .....	5
3.1. The "Basic" Traffic Engineering Formula .....	6
3.2. Broadband Subscriber Traffic Consumption .....	7
3.3. Network Capacity Modeling for Cable 10G Downstream.....	8
3.4. Network Capacity Planning for DOCSIS 4.0 ESD .....	9
3.4.1. Cable 10G Service Tiers .....	9
3.4.2. DOCSIS 4.0 ESD – Tavg and Nsub targets .....	10
3.4.3. DOCSIS 4.0 ESD – Network Capacity Requirements .....	11
4. Simulation Model Overview.....	11
4.1. Cable Model – PSD and TCP .....	11
4.2. Amplifier Model.....	11
4.3. System Noise Performance .....	13
4.4. Amplifier Spacing – Typical and Stretch Networks .....	13
5. Conventional Wisdom Using Unity Gain Approach.....	14
6. Game Changer – DOCSIS Compensates When Unity Gain Lags .....	17
6.1. Estimating HP/Subscriber Distribution across N+X HFC Plant.....	20
6.2. DOCSIS Optimized Network Capacity Analysis .....	21
6.3. Impact of shorter drop cables (66'/100'/150').....	22
6.4. Impact of Super-stretched Links (56dB) .....	24
6.5. Downstream ESD Capacity for various US/DS splits .....	26
6.6. Comparing 108-1218MHz FDX vs. 606-1794MHz ESD DS Capacity.....	29
6.7. Improving on the DOCSIS Weighted Average Capacity.....	29
6.8. Handling the Corner Cases – Meeting Tmax Burst QoE .....	31
7. Cost and Logistic Analysis .....	33
8. Conclusion.....	36
Acknowledgements .....	37
Bibliography & References.....	38
Abbreviations .....	39

## List of Figures

Title	Page Number
Figure 1 – A big picture view of a Hybrid Fiber Coaxial (HFC) network .....	4
Figure 2 – Illustration of the "Unity Gain" concept, applied to an RF amplifier cascade.....	5
Figure 3 – Downstream Average Bandwidth per Subscriber through Jan '20.....	7
Figure 4 – Upstream Average Bandwidth per Subscriber through Jan '20 .....	7
Figure 5 – 1218/204 MHz System – Subs per SG, DOCSIS Spectrum Needs.....	8
Figure 6 – 1218/204 MHz System – DOCSIS Usage: Tmax, Tavg, IP Video .....	9
Figure 7 – PSD Profile Showing Tilt/Drop/Tilt Characteristic.....	12
Figure 8 – Inter-amplifier Span Attenuation & Amp Gain (With 44 dB Gain limit) in a Stretched Plant.....	12
Figure 9 – Amplifier Spacing – Single Output Line Extenders (LE) .....	13

Figure 10 – Amplifier Spacing – Multi-Output Bridgers (MB).....	14
Figure 11 – End-of-Line Throughput - 108-1218 MHz, 150' Drop cables.....	15
Figure 12 – Stretch 108-1218 MHz Plant Bit-loading – Tap 5 for N+2 and N+5 .....	15
Figure 13 – End-of-Line Throughput - 492-1794 MHz, 150' Drop cables.....	16
Figure 14 – Typical 492-1794 MHz Plant Bit-loading – Tap 5 for N+2 and N+5 .....	16
Figure 15 – Stretch 492-1794 MHz Plant Bit-loading – Tap 5 for N+2 and N+5 .....	17
Figure 16 – 108-1218 MHz Net Throughput vs. Tap Position in N+6 Plant, 150' Drop cables .....	18
Figure 17 – 492-1794 MHz Throughput vs. Tap Position, N+6 Typical Plant, 150' Drops .....	18
Figure 18 – 492-1794 MHz Throughput vs. Tap Position, N+6 Stretch Plant, 150' Drops .....	19
Figure 19 – End-of-Line vs. Weighted Avg Throughput - 108-1218 MHz, 150' Drop cables.....	21
Figure 20 – End-of-Line vs. Weighted Avg Throughput - 492-1794 MHz, 150' Drop cables.....	22
Figure 21 – Net Throughput vs Tap Position in N+6 Typical Plant, 66/100/150 ft drop cable.....	23
Figure 22 – Net Throughput vs Tap Position in N+6 Stretch Plant, 66/100/150 ft drop cable.....	23
Figure 23 – 492-1794 Typical Plant with 1 or 2 Super-stretch 56 dB Links.....	25
Figure 24 –Weighted Average DS Capacity with 1 or 2 Super-stretch 56 dB Links.....	26
Figure 25 – Net Throughput vs. Tap Position in N+6 Typical Plant, different US/DS splits .....	27
Figure 26 – Weighted Average DS Capacity for different US/DS splits .....	28
Figure 27 – Weighted Avg Throughput - 108-1218 MHz vs. 606-1794 MHz, 150' Drop cables .....	29
Figure 28 – Example Bit-loading for Near and Far Homes .....	30
Figure 29 – Example Channel Capacities for various Tap Positions.....	30
Figure 30 – I-CCAP N+4 HFC network with 42/54 MHz sub-split .....	33
Figure 31 – CAPEX estimates for 1218/204 MHz N+4 (both without and with new tap faceplates), 1218 MHz FDX N+0, 1794 MHz ESD N+4 upgrades .....	34
Figure 32 – DAA N+0 HFC Network Model for 1.2 GHz FDX Upgrade.....	34
Figure 33 – DAA N+4 HFC Network Model for 1.8 GHz ESD upgrade .....	34
Figure 34 – Monte-Carlo analysis of “premium” for 1.8 GHz ESD case .....	35
Figure 35 – EoL vs. Weighted Avg Throughput – Typical and Stretch Plants.....	37

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Summary of SLA Options for 10G PON & 1218 MHz Plants .....	10
Table 2 – Example RF Amps and Homes Passed per Node for N+X .....	21
Table 3 – Summary of SLA Options for 1794 MHz Plants.....	28
Table 4 – Estimate of Labor force required to replace amps and taps in 5 years .....	36

# 1. Introduction

As of April 2020, all the 1,000+ pages of the Data Over Cable Service Interface Specification (DOCSIS) 4.0 PHY and MULPI specs, combined, are out and about. That's step one in making the 1,794 MHz top frequency in the downstream, combined with up to 684 MHz for the top frequency in the upstream, a new reality for the hybrid fiber-coax (HFC) plants. As part of the greater 10G initiative, this is what will be needed to support the many Gbps in both the upstream and the downstream for nodes with amplifier cascades (a.k.a. N+X). The multiple system operators' (MSOs') successful business model and existing HFC broadband networks are our starting point. All that's needed is a simple upgrade to get from here to there, easy peasy!


Not so fast?! Well, OK, there are some “minor details” to sort in the process of getting there: Where is the Goldilocks zone – not too little but not too much either – for the RF amps gain and power? There are many items to consider.

With a comprehensive network model of a node + 6 cascade, should the “unity gain” concept be extended all the way up to 1.8 GHz? Once modeled, what would an optimal power spectral density (PSD) distribution across the forward spectrum look like? What happens as the distance between the amps increases? Does the tap position on a link make a difference? What is the impact of the drop cables? Should we continue to use end-of-line (EoL) throughput as our capacity benchmark?

There are many, many questions and this paper will start to look at some of the considerations with rolling out 1794 MHz HFC plants. All of this is distilled into conclusions and best practices guidance – to help make 10G and the DOCSIS 4.0 networks a reality for the time to come!

## 2. Unity Gain Overview

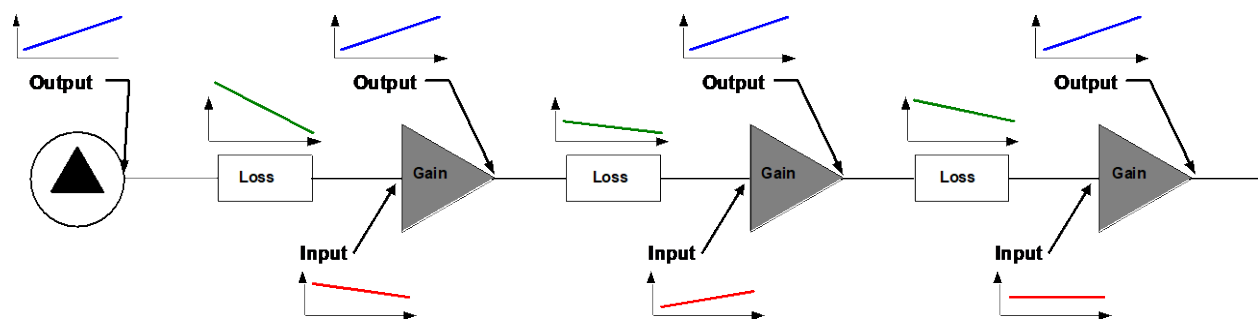
Figure 1 is a big picture view of a Hybrid Fiber Coaxial (HFC) network. It shows a “40,000-foot view” of an HFC plant. While the plant and its architecture kept evolving – for example, a head end to node fiber link took the place of a much longer cascade of trunk amps – it is also amazing how many things remain the same: fiber nodes feeding RF amplifiers, amplifiers feeding taps, and taps, via “drops” delivering signals to ever-important subscribers.



**Figure 1 – A big picture view of a Hybrid Fiber Coaxial (HFC) network**

Another HFC aspect that has remained the same is that of “unity gain” – a concept of setting every amplifier output signal to the same shape and magnitude, no matter if the amp is first or last in the

cascade nor how “lossy” the span in front of the said amplifier was. This consistency of output is shown via blue “amplifier output” lines in Figure 2 for the downstream direction.



**Figure 2 – Illustration of the "Unity Gain" concept, applied to an RF amplifier cascade**

In this example of a three-amplifier stage cascade, each output stage (in blue) is the same. However, the coaxial cable and tap loss preceding each of the amplifiers is of a different value, as illustrated via green “loss over frequency” lines. These various length cable spans, driven by the same previous-amplifier output, will produce various inputs, shown via red “amplifier input” lines. It is the gain and slope of every amplifier that is tuned, traditionally via selection of proper attenuator, cable simulator, and/or cable equalizer, to make every amplifier output the same. One goal of the tuning is to set each amplifier’s input stage signal as flat as possible to minimize noise figure contribution, the other is to up-tilt the final output, to minimize non-linear distortion effects. Nevertheless, the noise effects and the non-linear distortions still add up as the cascade length increases, despite keeping outputs of each amplifier the same.

Another rationale for unity gain is to provide a common input power and modulation order (e.g. 256-QAM) for every cable consumer premises equipment (CPE) including legacy video set top boxes (STB) as well as DOCSIS 2.0 &/or 3.0 cable modems (CM). This holds no matter what channel frequency is being used or whether the CPE device is next to the node or at the end of line (EoL).

The span length between the amps, and how much of the corresponding amplifier gain is needed to achieve “unity gain” has been the cable engineers’ focus for many years. Many HFC plants were originally designed with amp spacing to accommodate lower frequencies, e.g. 450/550/750 MHz. Over time, these were then “pushed” to higher frequencies, e.g. 870/1002/1218 MHz. Due to much higher cable losses at higher frequencies, this resulted in amplifiers needing significantly higher output gains to maintain the unity gain in the system. If the cable losses were more than the increased amplifier output gains, then the HFC plant needed to be “re-spaced” where amplifier locations were moved, something that is highly undesirable. This issue now gets exaggerated with pushing the HFC to 1.8 GHz.

As a point of reference, consider 750 MHz HFC plant that many “first world” countries built quite a few years ago. At 750 MHz, these spans amount to ~20 dB attenuation following trunk amps, ~40 dB attenuation following multi-port bridgers and ~30 dB attenuation following single output line extenders.

### 3. Network Capacity Planning for Cable 10G

The network capacity analysis in [ULM\_2019] provides an insight into the capacity requirements needed to support the cable 10G initiative. It first looks at the traffic engineering needed for a common 10G network using both PON and cable systems. Then a closer look is taken at the spectrum planning for an HFC system.

The CommScope (formerly ARRIS) team has been providing industry leading research in traffic engineering for many years which was recently highlighted in [ULM\_2019]. Some additional references of note include [CLO\_2014], [EMM\_2014], [ULM\_2014], [CLO\_2016], [ULM\_2016], [ULM\_2017] and [CLO\_2017].

### 3.1. The “Basic” Traffic Engineering Formula

Previously, [CLO\_2014] introduced traffic engineering and quality of experience (QoE) for broadband networks. From there, the paper went on to develop a relatively simple traffic engineering formula for service groups that is easy to understand and useful for demonstrating basic network capacity components.

The “Basic” formula shown below is a simple two-term equation. The first term ( $N_{sub} \cdot T_{avg}$ ) allocates bandwidth capacity to ensure that the aggregate average bandwidth generated by the  $N_{sub}$  subscribers can be adequately carried by the service group’s bandwidth capacity. The first term is viewed as the “DC component” of traffic that tends to exist as a continuous flow of traffic during the peak busy period.

#### The “2014” Traffic Engineering Formula (Based on $T_{max\_max}$ ):

$$C \geq (N_{sub} \cdot T_{avg}) + (K \cdot T_{max\_max}) \quad (1)$$

where:

$C$  is the required bandwidth capacity for the service group

$N_{sub}$  is the total number of subscribers within the service group

$T_{avg}$  is the average bandwidth consumed by a subscriber during the busy hour

$K$  is the QoE constant (larger values of  $K$  yield higher QoE levels)...

where  $0 \leq K \leq \text{infinity}$ , but typically  $1.0 \leq K \leq 1.2$

$T_{max\_max}$  is the highest Service Tier (i.e.  $T_{max}$ ) offered by the MSO

There are obviously fluctuations that will occur (i.e. the “AC component” of traffic) which can force the instantaneous traffic levels to both fall below and rise above the DC traffic level. The second term ( $K \cdot T_{max\_max}$ ) is added to increase the probability that all subscribers, including those with the highest service tiers (i.e.  $T_{max}$  values), will experience good QoE levels for most of the fluctuations that go above the DC traffic level.

The second term in the formula ( $K \cdot T_{max\_max}$ ) has an adjustable parameter defined by the  $K$  value. This parameter allows the MSO to increase the  $K$  value and add bandwidth capacity headroom that helps provide better QoE to their subscribers within a service group. In addition, the entire second term is scaled to be proportional to the  $T_{max\_max}$  value, which is the maximum  $T_{max}$  value that is being offered to subscribers.

In previous papers [CLOONAN\_2013, EMM\_2014], found that a  $K$  value of  $\sim 1.0$  would yield acceptable and adequate QoE results. [CLOONAN\_2014] goes on to provide simulation results that showed a value between  $K=1.0$  and  $1.2$  would provide good QoE results for a service group of 250 subscribers. Larger service groups (SGs) would need even larger values of  $K$  while very small SGs might use a  $K$  value near or less than  $1.0$ .

### 3.2. Broadband Subscriber Traffic Consumption

CommScope/ARRIS has been monitoring subscriber usage for over a decade now from the same group of MSOs. The data from this set has been compared and aligns closely to many other MSOs globally.

Figure 3 shows the average subscriber downstream consumption, DS Tavg, during peak busy hours for several MSOs over a ten-year period. At the start of 2020, DS Tavg had surpassed the 2 Mbps barrier.




Figure 3 – Downstream Average Bandwidth per Subscriber through Jan '20




Figure 4 – Upstream Average Bandwidth per Subscriber through Jan '20

It turns out that the Tavg growth rate was higher at the start of this decade and has tailed off a bit in recent years. Over the last 3-4 years, this group of MSOs had an average downstream traffic growth that had been around 30%. Interestingly, the upstream traffic is growing at a significantly slower rate than the



downstream as shown in Figure 4. During the same ten-year period, the upstream Tavg generally grew at less than 20% compound annual growth rate (CAGR).


This MSO data provides a good indication of Tavg, at least before the Coronavirus bandwidth (BW) surge hit. Note that figures 3 and 4 are very generalized results that are averaged across millions of subscribers.

Over recent years, there has been a slowing in the downstream usage growth rate (i.e. Tavg) compared to the service tier growth rate (i.e. Tmax). This has several consequences including that the networks become more “bursty”. It also means that the overall utilization of the network is lower too. In this respect, it is important to try and maximize subscribers per service group (SG) in order to take advantage of statistical multiplexing and get better economics.


### 3.3. Network Capacity Modeling for Cable 10G Downstream

[ULM\_2019] used the CommScope network capacity modeling tools to see how a 1218/204 MHz HFC plant could support the cable 10G downstream requirements. This analysis will first be reviewed and then look at what changes for the extended spectrum DOCSIS 4.0 (ESD) scenarios.

The network capacity results from [ULM\_2019] show the potential capabilities for a 1218/204 MHz HFC plant. It begins with a 512-home passed (HP) service group with 256 subs (i.e. 50% penetration). The max downstream (DS) service tier starts at 1 Gbps and grows by 1 Gbps per year from 2022 until it finally reaches 8 Gbps DS SLA in the year 2028. For this case study, it is assumed the Tavg growth rate continues its gradual decline over the next decade. This will leave Tavg at ~15 Mbps by the end of the decade. If the Tavg growth rate does not decline, then these dates might get pulled in by 2-3 years.



**Figure 5 – 1218/204 MHz System – Subs per SG, DOCSIS Spectrum Needs**



**Figure 6 – 1218/204 MHz System – DOCSIS Usage: Tmax, Tavg, IP Video**

This 1218/204 MHz scenario shows that by 2027, all SGs need to be at 128 subs or fewer. The max subs per SG is shown in Figure 5. The figure also breaks out the downstream spectrum amount needed for both DOCSIS 3.0 SC-QAM and DOCSIS 3.1 OFDM. Note that by 2028 in this case study, 100% of DOCSIS cable modems have been converted to DOCSIS 3.1 enabling OFDM channels across the entire spectrum to maximize capacity.

The DOCSIS capacity usage is broken out in Figure 6. It shows the amount of capacity needed for both DOCSIS 3.0 and DOCSIS 3.1. As can be seen, the Tmax component dominates over time. The upper red line shows the combined 3.0 + 3.1 total capacity for the system.

Perhaps a key point of this case study was that a reasonably clean 512 HP node can be upgraded to 1218/204 MHz and support a service tier of 8 Gbps x 1.5 Gbps for the next decade. The only change needed will be a SG segmentation (i.e. upgrade node from 1x1 to 2x2) somewhere in the middle of the decade. There is no pressing near term need to push the HFC to very small (but inefficient!) SG sizes found in N+0 systems.

### 3.4. Network Capacity Planning for DOCSIS 4.0 ESD

The cable 10G initiative strives for a more symmetric network with multi-gigabit upstream service tiers. DOCSIS 4.0 ESD has evolved as a way to offer even higher upstream tiers while remaining in an N+X environment. Thus, ESD also is trying to avoid the costs of an N+0 upgrade such as that needed with DOCSIS 4.0 FDX.

#### 3.4.1. Cable 10G Service Tiers

As shown in [ULM\_2019], the subscriber is actually getting an 8 Gbps SLA in a 10G world after accounting for all the different overheads (e.g. PHY, MAC, IP layers). Table 1 summarizes the various

options from that study and their respective downstream (DS) and upstream (US) SLAs that service providers can consider offering. Because capacity in an HFC system can vary quite a bit based on many variables, the offered SLAs are actually a range of values.

**Table 1 – Summary of SLA Options for 10G PON & 1218 MHz Plants**

<b>10G PON Options</b>	<b>DS SLA (Gbps)</b>	<b>US SLA (Gbps)</b>
10G/1G EPON	8	0.8
10G/10G EPON	8	8
XG-PON	8	2
XGS-PON, NG-PON2 (single wavelength)	8	8
<b>10G HFC Options</b>	<b>DS SLA (Gbps)</b>	<b>US SLA (Gbps)</b>
1218/85 MHz	8 – 10	0.4 – 0.5
1218/204 MHz	6 – 8	1.0 – 1.5
1218/300 MHz	5 – 7	2.0 – 2.5
1218/396 MHz	4 – 6	2.5 – 3.0
1218/85 MHz + 108-684 MHz FDX/Soft-FDX	8 – 10	5 – 6

Perhaps the key motivating driver for DOCSIS 4.0 FDX and ESD is the ability to offer more symmetric multi-gigabit upstream tiers. The minimum goal for ESD would be to at least match the DS tiers for a 1218/204 MHz plant while offering substantially higher upstream tiers. Ideally, ESD would be able to match the DS tiers capabilities of 10G PON and 4.0 FDX systems.

The table above shows upstream splits up to 396 MHz. The DOCSIS 4.0 working group has also added a 492 MHz split and 684 MHz split as additional options.

### **3.4.2. DOCSIS 4.0 ESD – Tavg and Nsub targets**

The [ULM\_2019] study assumed that the Tavg growth rates continue to decline and estimated that Tavg would reach 15 Mbps per sub by 2030. Given the additional investments needed for ESD, a Tavg of 20-40 Mbps per sub will be considered to make sure there is additional lifetime for this investment.

Since ESD is targeted at N+X environments, our analysis will assume at least 125 subscribers per SG, which was the ending point in the 2019 study.

### **3.4.3. DOCSIS 4.0 ESD – Network Capacity Requirements**

The 2019 study showed that ~9 Gbps of DS capacity is required for 125 subs @  $T_{avg}=15$  Mbps. If  $T_{avg}$  is bumped up to 25 Mbps, then total system capacity pushes above 10 Gbps. Pushing  $T_{avg}$  up to 40 Mbps increases the required system capacity to around 12 Gbps. These are key targets that will be used in our following analysis.

## **4. Simulation Model Overview**

Our network performance estimates were done using a MATLAB simulation model. The model simulated the effects of attenuation in the hardline coaxial cable, taps, drop cable and home entry port, and amplification in the intermediate amplifiers, to calculate the received power levels at modems at different locations. The effects of noise and distortion accumulation in the cascade were also simulated to calculate the effective SNR at the modem. The results were compared with DOCSIS 4.0 received power and SNR thresholds to estimate the available modulation and bit-loading at each frequency, yielding the total gross throughput. Physical layer overheads such as cyclic prefix, FEC (LDPC), guard bands, pilot and PLC tones were subtracted from the integrated bit-loading curve (gross throughput) to calculate the net system throughput. The model used 78% efficiency for the PHY layer.

### **4.1. Cable Model – PSD and TCP**


The DS power spectral density (PSD) at the node output was represented as a tilt/drop/tilt (a.k.a. lightning bolt) in order to achieve a sensible total composite power (TCP) (nominally 70 dBmV based on 4.0 working group discussions, but this parameter could be varied if needed) while providing adequate levels for legacy SC-QAM services at lower frequencies (up to 860 MHz). The PSD is shown in Figure 7.

Cable attenuation, for both hardline and drop cables, was calculated as a function of frequency and cable segment length from tabulated data from cable manufacturers. A large database of tap response curves from several tap families, taken from a combination of manufacturers' data and our own measurements, was available in the model. Note that no tap equalizers (or "cable simulators") were present in the model; all available RF power was used to optimize the bit-loading from each tap.


### **4.2. Amplifier Model**

The amplifiers were modelled as non-ideal near-unity gain devices. The gain of each amplifier was tuned to approximately match the attenuation of the preceding cable span, subject to a configurable maximum gain. This approximate match used a linear-tilted gain over most of the spectrum, producing a small gain deviation; over longer cascades, this produced a slightly U-shaped end-of-line PSD, in contrast to a flat received PSD which would occur if perfect unity-gain amplification was possible.

The gain limit became significant for longer cable runs that is referred to as stretched plant. At higher frequencies, the inter-amplifier span attenuation sometimes exceeded the assumed maximum gain, resulting in progressively lower PSD at these frequencies after each stage in the cascade. Hence the stretched plant did not quite achieve unity gain at 1800 MHz. Figure 8 shows an example of the inter-amplifier attenuation and amplifier gain for a stretched plant.



**Figure 7 – PSD Profile Showing Tilt/Drop/Tilt Characteristic**



**Figure 8 – Inter-amplifier Span Attenuation & Amp Gain (With 44 dB Gain limit) in a Stretched Plant**

### 4.3. System Noise Performance


The noise performance of the system was modelled using an estimate of the node's output SNR, and noise figures for the amplifiers in the cascade, and for the cable modem. The various noise components were modelled through the successive attenuation and amplification stages of the cascade to produce a received noise estimate. Distortion was treated in a similar way; distortion at the node output and each amplifier output was modelled as a third-order function of TCP, and coherence effects were also calculated.

### 4.4. Amplifier Spacing – Typical and Stretch Networks


For our network simulations, it was important to model amplifier spacings that were representative of real-world data based on customer's amplifier spacing data. A cumulative distribution function (CDF) of some sample amplifier spacing for single output line extenders (LE) is shown in Figure 9, while the CDF for multi-port bridgers (MB) is shown in Figure 10. Line extenders often account for more than 60% the total number of amplifiers. This data is courtesy of Cox Communications.

Most of the simulations used either “typical plant” or “stretch plant” parameters for plant length. “Typical plant” parameters are chosen to include a majority of inter-amplifier span attenuations (for line-extender (LE) spacings). A total attenuation of 30 dB at 1 GHz and 40 dB at 1794 MHz is used, which is an attenuation value greater than that of ~65% of line-extender to line-extender spacings according to Figure 9. This attenuation is simulated with 6 segments of 175' of P3 625 hard-line and 5 taps, a total of 1050' between amplifiers. Note that 850' of P3 500 would give similar results.

Unless otherwise specified, each simulation used the same parameters for the node to first amplifier span and for the last amplifier to end tap as for the inter-amplifier spans.



**Figure 9 – Amplifier Spacing – Single Output Line Extenders (LE)**



**Figure 10 – Amplifier Spacing – Multi-Output Bridgers (MB)**


“Stretch plant” parameters are chosen to exceed most inter-amplifier span attenuations; the corresponding attenuations are 35 dB at 1 GHz and 47 dB at 1794 MHz. This is greater than ~97% of LE-to-LE spacings, and two-thirds of bridger-to-amplifier spacings according to Figure 10. Overall, this would cover more than 85% of the amplifier links. This attenuation is simulated with 6 segments of 215’ of P3 625 hard-line and 5 taps, a total of 1290’ between amplifiers. Note that 1050’ of P3 500 would give similar results.

For both typical and stretch plants, the baseline case used 150’ of RG-6 drop cable along with a 3.5 dB loss inside the home (e.g. splitter at point of entry or 20’ of RG-6 inside the home). Note that cable losses from 150’ of RG-6 drop cable would be roughly equivalent to 210’ of RG-11 drop cable.

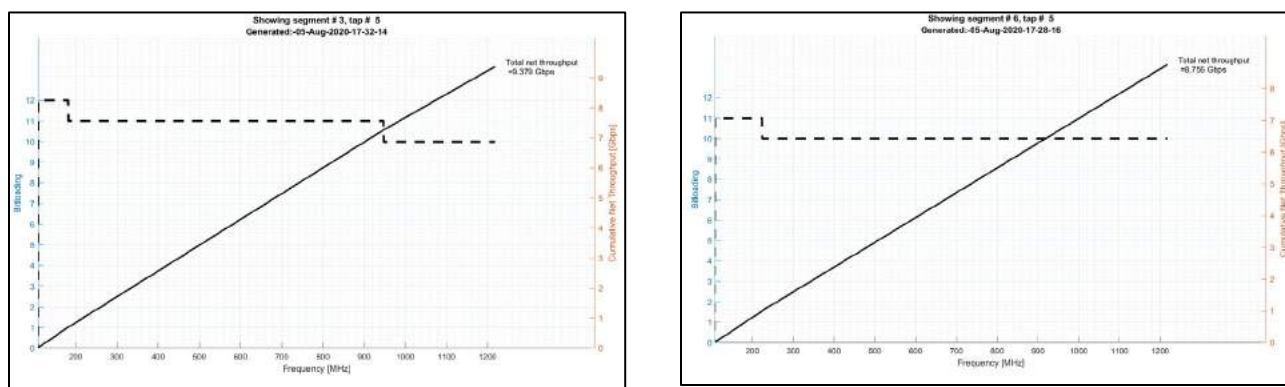
## 5. Conventional Wisdom Using Unity Gain Approach

As discussed previously, unity gain helped maintain a constant QAM-modulation for legacy STB and DOCSIS 2.0/3.0 cable modems. System performance was defined by the lowest common denominator. The end of line (EoL) provided a key performance monitoring spot to determine whether every CPE device could successfully receive the 256-QAM downstream channel. But future systems will be migrating to DOCSIS 3.1 and 4.0 cable modems. Our network simulations calculate the maximum downstream network capacity achievable using 3.1/4.0 OFDM channels with variable-bit loading per sub-carrier with up to 4096-QAM modulation.

For our network simulations, the EoL performance is first checked for both 1218 and 1794 MHz systems to see how well unity gain performed. Figure 11 looks at the network capacity for a 108-1218 MHz downstream for various amplifier cascade lengths, from Node+0 up to Node+6. This was done for both a typical plant and a stretch plant conditions.



**Figure 11 – End-of-Line Throughput - 108-1218 MHz, 150' Drop cables**



**Figure 12 – Stretch 108-1218 MHz Plant Bit-loading – Tap 5 for N+2 and N+5**


The 1218 MHz system maintains unity gain for both the typical and stretch plants. The network capacity degrades slightly as the amplifier cascades increases due to the accumulated noise and distortion from each additional amp stage. The additional cable loss in the stretch plant degrades network capacity slightly from the typical plant, but each amplifier stage is still able to maintain its unity gain.

Figure 12 shows the bit-loading for the 5<sup>th</sup> tap in a N+2 and N+5 stretch 108-1218 MHz system. As can be seen, the bit-loading remains fairly constant over the entire spectrum gain. This is what one might expect with unity gain.

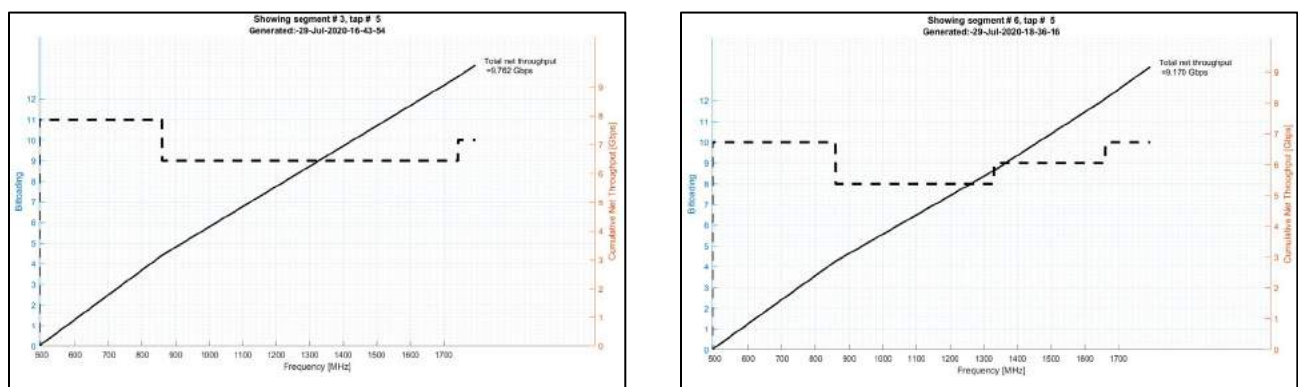


Figure 13 now shows the EoL throughput for a 492-1794 MHz downstream system with cascade lengths varying from Node+0 to Node+6. This has 192 MHz more total downstream spectrum than the 108-1218 MHz system, but it is spectrum at higher frequencies with higher losses. Note that the downstream network capacity of the 492-1794 MHz typical plant in Figure 13 is very close to the 108-1218 MHz typical plant in Figure 12. However, the EoL throughput for the 1.8 GHz stretch plant in Figure 13 is substantial lower than the typical plant by around 20% or almost 2 Gbps. The stretch plant is also below the minimum 9 Gbps capacity target that was discussed in section 3.

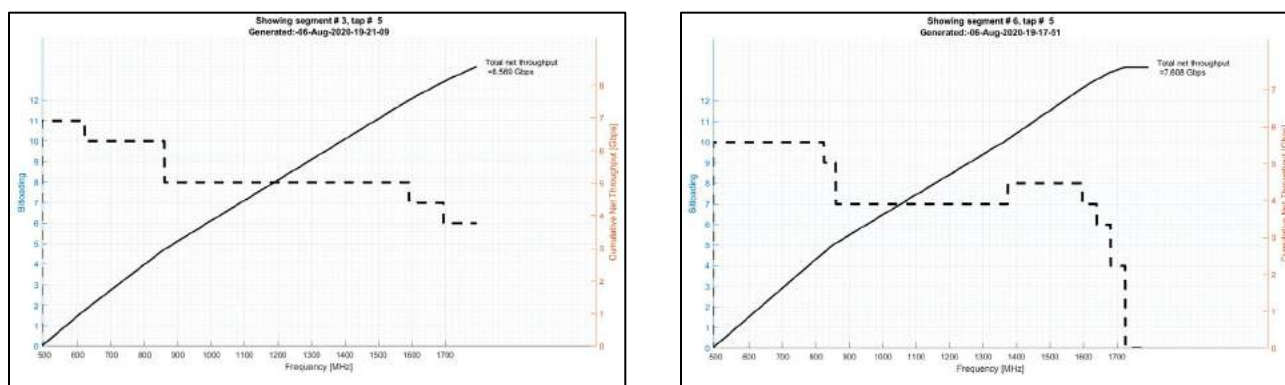
Figure 14 shows the bit-loading for the 5<sup>th</sup> tap in a N+2 and N+5 typical 492-1794 MHz system. Since unity gain is maintained, there is no drop-off in bit-loading at higher frequencies. In fact, there is a slight increase in the bit-loading at the higher frequencies due to some non-linear effects in the amplifier.



**Figure 13 – End-of-Line Throughput - 492-1794 MHz, 150' Drop cables**



**Figure 14 – Typical 492-1794 MHz Plant Bit-loading – Tap 5 for N+2 and N+5**



**Figure 15 – Stretch 492-1794 MHz Plant Bit-loading – Tap 5 for N+2 and N+5**

The stretch 492-1794 MHz system does not maintain unity gain. Figure 8 showed that the amplifier gain reaches its maximum output around 1600 MHz. Figure 15 shows the bit-loading for the 5<sup>th</sup> tap in a N+2 and N+5 stretch 492-1794 MHz system. It clearly shows how the bit-loading starts to drop-off above 1600 MHz. Note that the loss is relatively minimal after the 2<sup>nd</sup> amplifier output, only 1-2 orders of modulation. However, by the output of the 5<sup>th</sup> amplifier, the 5<sup>th</sup> tap is seeing a dramatic drop in bit-loading, losing roughly half the capacity in the 1602-1794 MHz OFDM channel.


Based on our earlier network capacity planning targets of needing 9-12 Gbps, the stretch 492-1794 MHz plant would not have acceptable system capacity using EoL throughput measure. Our conventional wisdom would dictate that these plants would need to have the amplifiers re-spaced or a mid-span amplifier added in order to maintain unity gain.

## 6. Game Changer – DOCSIS Compensates When Unity Gain Lags


When the DOCSIS 3.1 (D3.1) specification was written, the authors had a vision of a new DOCSIS PHY layer that squeezed every last ounce of capacity out of the cable plant. DOCSIS 4.0 inherits all these capabilities. Some of the most important D3.1 capabilities were wide OFDM channels with variable bit-loading along with the support of multiple profiles.

Previously, all modems would receive the identical downstream data stream. Therefore, the entire network had to operate for the lowest common denominator. Hence, the EoL capacity was a good metric for determining the plants capabilities. Each OFDM channel has a group of D3.1 profiles. Within that channel, modems are put into profile groups with similar modems and receive data at an optimal data rate. For example, modems in one profile could be receiving at 4096-QAM modulation while other modems in a different profile are receiving at 256-QAM modulation. This allows the DOCSIS system to optimize system capacity. This means that total system capacity depends on the capacity seen by every home on every tap in the system.

Figure 16 shows a 108-1218 MHz downstream example in a N+6 plant for both typical and stretch amplifier spacing. Network capacity is shown at every tap in the system. Modems further out on the cascade are seeing slightly lower performance and would be operating with D3.1 profiles that are using QAM modulations one or two steps below the best profile. In this example, the best profile would have 20% more capacity than the lowest profile.



**Figure 16 – 108-1218 MHz Net Throughput vs. Tap Position in N+6 Plant, 150' Drop cables**



**Figure 17 – 492-1794 MHz Throughput vs. Tap Position, N+6 Typical Plant, 150' Drops**


For 1.8 GHz plants, the cable losses at the higher frequencies will magnify the differences between the best profiles and the lowest profiles. Figure 17 helps give a glimpse into the wider discrepancy seen by modems across an entire typical plant. Capacity changes based on two factors: # of amplifiers passed; and # of taps from the amplifier output. Figure 18 shows the stretch plant that has even more of a swing. In this example, the typical plant has a 32% swing from the lowest to highest capacities while the stretch plant sees a 60% swing from low to high! It goes from ~7.5 Gbps at tap 5 on the 6<sup>th</sup> amplifier up to almost 12 Gbps at the 1<sup>st</sup> tap from the node.

Remember, the EoL capacities discussed earlier basically aligns with the 5<sup>th</sup> tap value after the last amplifier, so it is basically following the troughs on this chart (see EoL for each N+X in Figures 17 & 18). These wide variations are just waiting to be optimized by DOCSIS. Many of the modems will have capacities significantly above the EoL capacities.

The profiles are set up separately for each DOCSIS OFDM channel. The OFDM channels are up to 192 MHz wide. So, the 492-1794 MHz downstream will have up to 7 OFDM channels with each channel getting its own optimized profiles.

In addition to optimizing each channel, the CMTS scheduler does load balancing between the channels. And an intelligent scheduler can factor the profiles for each modem to determine which channels it should use to optimize total network capacity.

So, if EoL is not a good measure of the network capacity, what should be used? It turns out that this is an extremely complex answer based on many different variables. The following sections will explore a number of these variables and make some estimates of network capacity for a ‘reasonable’ system.



**Figure 18 – 492-1794 MHz Throughput vs. Tap Position, N+6 Stretch Plant, 150' Drops**

## 6.1. Estimating HP/Subscriber Distribution across N+X HFC Plant

As can be seen in Figures 17 & 18 above, the available capacity to any given modem is a function of where it is located on the plant relative to the fiber node. The HFC plant is a tree and branch topology that fans out for each additional level in the cascade. If one blindly assumed that each amplifier has 2 outputs feeding the next level of amplifiers, then a N+6 plant would have 64 amplifiers at the last stage in the cascade and a total of 126 amplifiers per node. If the outputs per amplifier is increased to 3 or 4, then it grows exponentially to where there could be 1000's of amplifiers from a given node.

But this is not reality. To understand what is out in the real world, the authors spoke with our customers and tapped into the knowledge of the in-house CommScope HFC design team. This in-house design team is perhaps the most experienced HFC design team anywhere, with a legacy going back 30+ years at some of our previous incarnations: ARRIS, C-Cor, Motorola, General Instrument, Philips.

The first thing we learned is that there is really no “typical” system. Real world HFC plants vary all over the map. So, we set out to define a “reasonable” scenario to get some baseline capacity estimates. From there, certain variables can then be changed to understand their impact on network capacity.

In general, amplifier spacing on average is relatively constant at 4-5 amplifiers per mile – especially if the networks considered include a mix of low density rural, medium density suburban and high-density urban areas. The homes passed (HP) density impacted how many homes might be off each tap (e.g. 1-2 HP per rural tap, 4-8 HP per urban tap). So, for our analysis, it is assumed that HP would be distributed evenly across all taps in the system.

As a starting point, a suburban build with 80-110 homes passed per mile (HP/mile) was chosen. From there, plants of various cascade lengths (i.e. N+0 to N+6) were analyzed. Table 2 shows an estimate of RF amps and homes passed based on input from our expert in-house HFC design team. This table shows that the number of RF amplifiers per node grows almost linearly with the cascade depth. It is NOT exponential! For each level in the cascade, another six to eight amplifiers are typically added.

Why is this??? It turns out there are several factors. First, the majority of amplifiers are single output line extenders (LE). These are amplifiers just cascading down a street without any branches. Next, multiple output bridgers (MB) would often have some outputs feeding side streets without connecting to another amplifier (i.e. a dead end). Finally, not all branches go out the entire cascade depth. In a N+6 plant for instance, many branches might terminate after 3, 4 or 5 amplifiers and not reach the sixth cascade level.

The homes passed per node also tended to follow the amplifier count with typically 16-20 HP per amp.

This knowledge then let us model how the homes passed (and hence subscribers) are spread across the HFC plant. After the first amplifier, a linear increase in HP per cascade level is used.


The next task was to understand how many taps are typically following each amplifier. As noted above, the amplifier spacing is fairly constant, and hence the number of taps were too. But there is some variation. The HFC design team might need to shorten or extend a leg based on the real-world geography. The vast majority of amplifiers have either 4 or 5 taps at its outputs. The number of fewer taps (e.g. 3) or more taps (e.g. 6) was trivial for this analysis. It turns out that percentage of amplifiers with 4 taps is roughly twice that of the amplifiers with 5 taps.

**Table 2 – Example RF Amps and Homes Passed per Node for N+X**


HFC Cascade Lengths	# of RF Amps per Node	# of Homes Passed per Node
Node + 0	0	40 - 60
Node + 1	4 – 6	80 - 150
Node + 2	10 – 12	150 - 240
Node + 3	15 – 20	240 - 360
Node + 4	22 – 28	360 - 480
Node + 5	28 – 35	480 - 600
Node + 6	35 - 45	600 - 720

## 6.2. DOCSIS Optimized Network Capacity Analysis

The first order of business was to look at the 108-1218 MHz plant shown in Figure 16 and see how much additional DOCSIS capacity might be available compared to the EoL measurements. Figure 19 shows the DOCSIS optimized weighted average capacity compared to the EoL capacity for both the typical and the stretch 1218 MHz plants.



**Figure 19 – End-of-Line vs. Weighted Avg Throughput - 108-1218 MHz, 150' Drop cables**



**Figure 20 – End-of-Line vs. Weighted Avg Throughput - 492-1794 MHz, 150' Drop cables**

Both typical and stretch plants saw very little gain for shorter cascades while N+6 plant saw up to 10% gains over the EoL calculations. So EoL appears to be reasonably accurate for the 1218 MHz plant.

The next scenario looked at the 492-1794 MHz plant previously shown in Figure 17 & 18 to see how much additional DOCSIS capacity might be available compared to the EoL measurements. Figure 20 shows the DOCSIS optimized weighted average capacity compared to the EoL capacity for both the typical and the stretch plants.

The difference between the DOCSIS optimized weighted average capacity and EoL capacity is significant. For the typical plant, the gains start around 10% for N+0 and increase to 15% for N+6. The stretch plant sees even larger gains, from 25% gains for N+0 plant to 35% for N+6 plant. The EoL measure is grossly underestimating the capacity of the stretch plant. The EoL measure showed that typical plants had 20% to 25% more capacity than the stretch plant. The DOCSIS optimized capacity results above shows that the gap between typical and stretch plant is significantly less, in the 5% range!


Perhaps one of the most significant conclusions from this study is that the stretch plant that was considered inadequate, is actually capable of supporting the network capacity requirements and does NOT need any extra help (e.g. re-spacing amps or adding mid-span amplifiers).

### 6.3. Impact of shorter drop cables (66'/100'/150')


Drop cables are less costly than trunk lines, but also do not have as good RF characteristics. This becomes a big deal as the HFC plant pushes to 1.8 GHz. Per our in-house HFC design expert, most drop cables out there today are RG-6 cables. Higher performance RG-11 drop cables are typically used for extremely long runs, e.g. 200' or longer.

Our expert felt that our 150' drop cable assumption in the previous results was on the high side. Figure 21 and 22 show the impact on total network capacity by varying the drop cable length from 66' to 100' to

150'. This was done for the typical plant in Figure 21 and the stretch plant in Figure 22, using RG-6 cable for all cases.



**Figure 21 – Net Throughput vs Tap Position in N+6 Typical Plant, 66/100/150 ft drop cable**



**Figure 22 – Net Throughput vs Tap Position in N+6 Stretch Plant, 66/100/150 ft drop cable**



In general, the short drop cable makes a significant difference, especially for the fourth and fifth taps. For the typical plant, there is less than a 3% difference on the first two taps after the amplifier for 100' compared to 150'. The 4<sup>th</sup> tap has gains in the 7-10% range while the 5<sup>th</sup> tap sees 7-15% improvements.

The stretch plant shows even wider spreads between 100' and 150'. The gains on the first two taps can be up to 4%. The 4<sup>th</sup> tap sees improvements up to the 8-12% range and the 5<sup>th</sup> tap sees 10-16% gains with the shorter 100' drops.

*Note that replacing a 150' RG-6 drop cable with RG-11 will result in improved performance that is very close to the 100' RG-6 drop cable.*

Migrating from 100' to even shorter 66' drop cables did not have as big effect, especially on the typical plant. The 66' drop cable did show some gains over the 100' drop on the 4<sup>th</sup> and 5<sup>th</sup> taps of the stretch plant. The largest gains being 7-10% after the first two amplifier and then shrinking gains with more amplifiers in the cascade.

#### **6.4. Impact of Super-stretched Links (56dB)**

As shown on Figure 10, a small percentage of multi-port bridger amplifiers have extremely large dB spacings that are >55 dB. Figure 23 shows the impact of inserting one or two of these super-stretched links with 56 dB spacing into a typical plant. For reference, 56 dB spacing between amplifiers could represent any one of the following scenarios:




1. 1560 ft of 0.500 cable
2. 1920 ft of 0.625 cable
3. 2640 ft of 0.875 cable
4. Or 1560 ft of 0.625 plus 5 taps

In general, it would be expected that these super stretched links are express feeder runs over a long distance without any taps. For this example, Figure 23 shows scenario 4 with 1560' of 0.625 cable plus 5 1.8 GHz taps to demonstrate the impacts of the distance on these taps.


For the first case, a super-stretch 56 dB link is inserted following the node. As can be seen, the total capacity seen at the output of the 1<sup>st</sup> amplifier drops by ~0.5 Gbps at the 1<sup>st</sup> tap and over 1 Gbps by the 5<sup>th</sup> tap relative to the typical plant. After each additional amplifier in the cascade, the delta with the typical plant shrinks as the noise introduced by each amplifier accumulates and starts to dominate.

For the second case, the super-stretch 56 dB link is inserted after the 3<sup>rd</sup> amplifier. The outputs after the node, 1<sup>st</sup> and 2<sup>nd</sup> amplifiers are identical to the typical plant. The outputs after the 4<sup>th</sup>, 5<sup>th</sup> and 6<sup>th</sup> amplifiers are very close to the first case in which the super-stretch link is at the node.

The final case shows a scenario where two super-stretch links are inserted, one after the node and one after the 3<sup>rd</sup> amplifier output. Adding the second super-stretch link shows a noticeable degradation to the total network capacity after the 4<sup>th</sup>, 5<sup>th</sup> and 6<sup>th</sup> amplifier outputs.



**Figure 23 – 492-1794 Typical Plant with 1 or 2 Super-stretch 56 dB Links**



**Figure 24 –Weighted Average DS Capacity with 1 or 2 Super-stretch 56 dB Links**

Figure 24 compares the DOCSIS optimized weighted average DS capacity for the three super-stretch cases compared to the typical plant. With the CMTS scheduler effectively averaging capacity across all the taps, this helps to minimize the impact of these super-stretched links. For Node+6 plant, adding one super-stretched link drops average capacity by <5% while adding two super-stretched links only has ~10% impact.

### 6.5. Downstream ESD Capacity for various US/DS splits

The DOCSIS 4.0 specification allows for many different upstream (US) options. For a static ESD system, the possible diplexer splits are:

- 85/108 MHz
- 204/258 MHz
- 300/372 MHz
- 396/492 MHz
- 492/606 MHz
- 684/834 MHz

The first number is the upper edge of the upstream spectrum. The second number is the starting edge of the downstream. The guard band is the region between these two and is not usable. Notice that the guard band grows larger as the US/DS split frequencies increase.


Our network capacity modeling research has shown time and again that downstream bandwidth capacity is the limiting factor in almost every analysis. So, it is extremely important to try and maximize the downstream bandwidth capacity.

The DOCSIS 4.0 specification also supports a dynamic FDX option. For FDX, the shared band starts at 108 MHz and can go up to 204, 300, 396, 492 or 684 MHz. The FDX downstream goes up to 1218 MHz. So downstream bursts can potentially leverage the full 108-1218 MHz.

Figure 25 shows the downstream capacity at each tap location for various static ESD US/DS splits for a N+6 typical plant and 150' drop cables. It also compares these to a 108-1218 MHz DS similar to that used by FDX (although full-duplex operation may add other impairments which are not factored into this example). Figure 26 then shows the DOCSIS optimized DS weighted average capacity for the various static ESD splits and the dynamic FDX for various cascade lengths from N+0 to N+6.

Because the changes between the US/DS splits are in the lower frequencies (i.e. below 684 MHz), the capacity differences remain constant. For each static ESD curve, there is a drop in capacity with increasing taps. This is the same for all static ESD scenarios, so they all have the same shape. Note that the 684/834 split has a noticeably larger gap than the other splits. This is a combination of two factors. First, it has the largest guard band. Second, its downstream is missing a lot of the lower frequencies which tend to have higher bit-loading as shown previously in Figures 14 and 15.


The 108-1218 MHz FDX DS spectrum is also shown for reference and appears to align closest with the 492/606 MHz static ESD DS capacity. The next section takes a closer look at that.



**Figure 25 – Net Throughput vs. Tap Position in N+6 Typical Plant, different US/DS splits**

Table 3 looks at some potential service tier combinations that could be considered for the various ESD split options. These SLAs are based on the basic traffic engineering formula discussed earlier in section 3. Remember that it is a function of the number of subscribers,  $N_{sub}$ , and the average peak period consumption per sub,  $T_{avg}$ . Both may change over time. For this table, the base assumption is  $N_{sub}=200$ ,  $DS\ T_{avg}=20\ Mbps$ ,  $US\ T_{avg}=1.2\ Mbps$  and  $K=1.0-1.2$ . Note that 100 subs @  $DS\ T_{avg}=40\ Mbps$  would be roughly equivalent.

By comparison, the SLA for 10G PON with 64 subs and same Tavg would be 6-7.5 Gbps for both the US + DS. The last row in Table 3 shows the theoretical best SLA's on a 1794 MHz plant. This uses dynamic soft-FDD where the CMTS scheduler switches between two diplexer settings (i.e. 85/108 MHz and 684/834 MHz) to time division multiplex the 108-684 MHz spectrum between US + DS.



**Figure 26 – Weighted Average DS Capacity for different US/DS splits**

**Table 3 – Summary of SLA Options for 1794 MHz Plants**


10G HFC Options	Top of US (MHz)	Start of DS (MHz)	DS SLA (Gbps)	US SLA (Gbps)
1794/204 MHz	204	258	7 – 9	1 – 1.5
1794/300 MHz	300	372	6 – 8	1.5 – 2
1794/396 MHz	396	492	5 – 7	2.5 – 3
1794/492 MHz	492	606	4 – 6	3 – 3.75
1794/684 MHz	684	834	3 – 4	4 – 5
1794/85 MHz with 108-684 MHz dynamic Soft-FDD	684	108	8 – 10	4 – 5

## 6.6. Comparing 108-1218MHz FDX vs. 606-1794MHz ESD DS Capacity

The DOCSIS optimized capacity for both the static 606-1794 MHz ESD and the 108-1218 MHz dynamic FDX DS capacities are compared in Figure 27. This gives a better handle on how these two stacks up against each other. Note – FDX in N+X plant would require FDX-capable amplifiers which are still under investigation.

For typical plant conditions, the DOCSIS optimized weighted average capacity for the 606-1794 MHz ESD is almost identical to the 108-1218 MHz FDX system.

On a stretch plant, the 606-1794 MHz ESD takes a hit in the higher frequencies, especially above 1.5 GHz. This reduces its DS capacity by about 5% compared to the 108-1218 MHz DS in a stretch plant.




**Figure 27 – Weighted Avg Throughput - 108-1218 MHz vs. 606-1794 MHz, 150' Drop cables**

## 6.7. Improving on the DOCSIS Weighted Average Capacity


The DOCSIS weighted average capacity in the earlier analysis's assumes that the subscriber's data usage is relatively evenly distributed across the HFC plant AND the subscriber's data usage is evenly distributed across all the available OFDM channels. However, each channel has its own unique set of profiles that can vary significantly from channel to channel.

Figure 28 below looks at the bit loading for two extreme cases in the 492-1794 MHz plant. The first case is a home on the 1<sup>st</sup> tap after the 1<sup>st</sup> amplifier for a typical plant amp spacing. The second case is a home on the 5<sup>th</sup> tap after the 6<sup>th</sup> amplifier on a stretch plant.



**Figure 28 – Example Bit-loading for Near and Far Homes**

As can be seen, the modem at the far home has vastly lower bit-loading capacity in the upper frequencies, and in particular, the top OFDM channels, i.e. 1602-1794 MHz. If the far modem has a 5 Gbps burst, the CMTS could schedule it in the 492-1218 MHz OFDM channels. Note that the near modem still has ~5 Gbps available to it in the top 3 OFDM channels from 1218-1794 MHz. So, both modems could be bursting to 5 Gbps simultaneously even though the EoL throughput for this plant is less than 7.5 Gbps. And the total burst rate is higher than the average of the two bit-load maps put together.



**Figure 29 – Example Channel Capacities for various Tap Positions**

Figure 29 shows how these bit-loadings map to the individual channel capacity for the 1<sup>st</sup>, 3<sup>rd</sup> and 5<sup>th</sup> tap on every amplifier output stage. As can be seen in the above channel capacities, the top OFDM channel for the 5<sup>th</sup> tap is practically unusable. For the other high frequency OFDM channels (i.e. 1026-1602 MHz) the 1<sup>st</sup> and 3<sup>rd</sup> taps can be 30% to 50% higher capacities than the 5<sup>th</sup> tap location. Overall, the 1<sup>st</sup> tap position is getting reasonably good capacity across all the channels, even at the higher frequencies.

For over a decade, the CMTS schedulers have been load balancing across multiple DOCSIS 3.0 bonded channels. And they have been doing a great job at it. In the above scenario, the CMTS scheduler can tend to put data from the further taps in the lower frequency OFDM channels while the nearer taps utilize the relatively empty upper frequency OFDM channels. This means the CMTS can achieve even higher capacities than those shown with the weighted average capacity analysis above.

## 6.8. Handling the Corner Cases – Meeting Tmax Burst QoE

The DOCSIS weighted average capacity and the intelligent CMTS scheduler improvements addresses the total system capacity available to the scheduler. However, as seen previously, the capacity seen by any individual home can vary dramatically. The CommScope basic traffic engineering formula discussed in Section 3 can help determine the capacity needed for each individual home so it can obtain its appropriate Quality of Experience (QoE).

Consider a worst-case scenario where an operator wants to support a 7.5 Gbps DS service tier across their entire footprint, including 10G PON FTTH and DOCSIS 4.0 ESD 492/1794 MHz HFC plant. As seen in Figure 18, some home locations might only see total available capacity around 7.5 Gbps. That means that whenever that home wants to burst to its full 7.5G Tmax, it will need to use every ounce of spectrum available to it. That would leave practically nothing left for the remaining subs (i.e.  $N_{sub} \cdot T_{avg}$  component from the formula).

In reality, the CMTS would not let these other subs starve, but will tend to allocate higher capacity lower spectrum to the bursting modem and others near the EoL. Meanwhile modems closer to the node and amplifier outputs will be using their higher capacity profiles in the higher frequency OFDM channels. So, the bursting modem might only get 5-6 Gbps out of its 7.5G tier during peak busy times.

So, how probable is this worst-case scenario? There are several factors that must all align for this to happen. These factors and the individual probability of each include:

- Subscriber takes the premium 7.5G DS top billboard tier [~1% to ~5%]
- Subscriber home is on the 5<sup>th</sup> tap [20%, i.e. 1 out of 5]
- The probability that this link even has 5+ taps [~30%]
  - Majority of amplifier outputs have 4 taps or fewer
- The tap is after the 5<sup>th</sup> or 6<sup>th</sup> amplifier [~25% of amplifier outputs on N+5/N+6 plant]
- % of MSO plants that are even N+5 or higher [1% to 20%??]
- The amplifier output is being stretched beyond unity gain [~5% to 15%]
- The RG-6 drop cable is >125' [~25% to 50%]

On N+5/N+6 plants, only 1 out of every ~200 to ~1,000 amplifier outputs might fall into this category. While it varies greatly by MSO, N+5 and N+6 plants are also becoming a smaller and smaller percentage of their total HFC plants. Of these potential problem amplifier links, it is only a problem if there is a subscriber on the 5<sup>th</sup> tap that takes the 7.5G DS top billboard tier. For all these stars to align, it looks like a probability on the order of four or five 9's that this won't happen.



It may be the end of this decade before the 7.5G DS tier is being offered and this worst-case scenario can potentially kick in. So, what options are available to the operator over this time window to handle this potentially rare event? Here are some possibilities:

1. MSO accepts slight QoE degradation (i.e. 7.5G drops to 5-6G during peak busy hour for these very infrequent customers)
2. MSO only offers 4G-5G DS service tier to this customer
3. Pull fiber to that subscriber's home and switch them to 10G PON
4. Reduce the N+5/6 cascade depth
5. Add mid-span amplifiers to boost higher frequency signals
6. Replace the RG-6 drop cable with RG-11

Option 1 might be very acceptable, but potentially dependent on your regulatory environment. The customer perceived QoE between 5 Gbps bursts and 7.5 Gbps burst capacity may not be perceptible. And this is only happening for rare times when the consumers bursts to the max and for a very small percent of the customer population. If the operator is still uncomfortable with this, they can choose option 2 and only provide the consumer with a slightly reduced T<sub>max</sub> (e.g. 5G instead of 7.5G). Note, this is something that the DSL world has had to always deal with, but this is not nearly as dramatic. In older DSL technology, customer capacity might have varied from 2 Mbps to 25 Mbps.

If the operator wants to correct the situation and provide additional capacity where needed, then options 3-6 would need to be considered over the next 5-8 years. An economic analysis of some of these choices is in the next section. Since this scenario does not need to get resolved until the end of the decade, the operator has plenty of time to correct the issue before it happens.

Everyone agrees that the long-term strategy is to eventually get to FTTH. However, many consider this a multiple decade transition to get to the point where fiber is pulled down every street. Option 3 jumps directly to FTTH as the solution. However, since this is a N+5/N+6 plant, chances are that the fiber is not nearly as close to the home as it might be in N+0/N+1 plant. It might be over a mile away from the customer. So, option 3 is expected to be the most expensive.


With the long-term strategy to continually push fiber deeper, then option 4 above is making a step in that direction. If these problem amplifier spans can be identified today (e.g. N+5/N+6 with stretched outputs), then the operator can spend the next decade gradually attending to these and pushing fiber deeper. Grey optics aggregator (GOA) / grey optics terminator (GOT) is a cost-effective fiber deep architecture that could be used to reduce the cascade length, without necessarily going all the way to N+0. The GOA functionality is placed at the parent node location. The downstream optical signal is split and passed on to the GOT node using short distance optics. The GOA also aggregates all of the GOT return signals. The GOT nodes are transparent to the head end and does not require any additional head end optics.

A low cost GOT fiber node can replace one of the multi-bridger amplifiers on the path to these problem spans to reduce the cascade length. If an operator has a 20 year plan to convert to a fiber deep network (e.g. 5% of plant per year), they can target these problem spans now to make sure issues with these are corrected before the end of the decade. Later in time, the GOT node can be upgraded to a full-fledged fiber node as needed.

Options 5 and 6 both fix the potential capacity issue but are not in alignment with the strategic direction of pushing fiber deeper towards eventual FTTH. Option 6, pulling a new drop cable, could be aligned with this strategy if it is a "siamese" cable that contains both fiber and RG-11 drops. Option 5, adding a mid-span amplifier, might make sense if the logistics do not allow option 4 to pull fiber deeper to reduce cascade lengths.

## 7. Cost and Logistic Analysis

Our starting point for economic analysis discussion is a 750 or 860 MHz HFC network case study depicted in Figure 30. The fiber node area, serving 400 HP, is also a single Integrated Converged Cable Access Platform (I-CCAP) service group. It is a 42/54 MHz “sub-split” system with a total of 20 RF amplifiers, 10 each of bridger and line extender type. There is a total of 100 taps, on 25,000 feet of hardline coax plant, for an average density of 84 HP/mile and ~4.2 amps/mile. The longest RF amp cascade length is N+4.




**Figure 30 – I-CCAP N+4 HFC network with 42/54 MHz sub-split**

The cost of upgrading this network to 1218 MHz with a 204/258 MHz high-split is analyzed by (a) replacing “ePack” node and amps modules and (b) adding appropriate I-CCAP license to augment both DS and US data capacity. The left-most column of Figure 31 depicts this network upgrade, along with the percent breakout of various elements. Furthermore, this 1218/204 MHz upgrade is normalized to 100%, in order to be the baseline to compare to the other cases, namely 1.2 GHz FDX N+0 and 1.8 GHz ESD N+4. The 2<sup>nd</sup> column shows the 1218/204 MHz upgrade where the tap faceplates are also upgraded to 1.2 GHz. This adds ~25% premium on top of the base case.


The network shown in Figure 30 is applicable to the 1.2 GHz high-split case. For the other two cases, however, networks of Figure 32 and Figure 33 are the respective representations. FDX requires a N+0 “fiber-deep” upgrade. The original fiber node plus 20 amplifiers end up being replaced by 4 fiber nodes with ~50% of the original hardline coax over lashed by new fiber. For the 1.8 GHz ESD N+4 case, whole-station amplifiers, as well as all 100 tap housings are replaced.

As shown in Figure 31, the FDX N+0 upgrade comes to 454% of the “base case”, that is, the 1.2 GHz high-split upgrade. The fiber material plus labor costs really dominate for the N+0 upgrade. Our analysis assumed 80% aerial plant. This component could be much higher if the percentage of underground plant increases over 20%. The 1.8 GHz ESD upgrade comes to 214% of the “base case”, or a 114% premium.

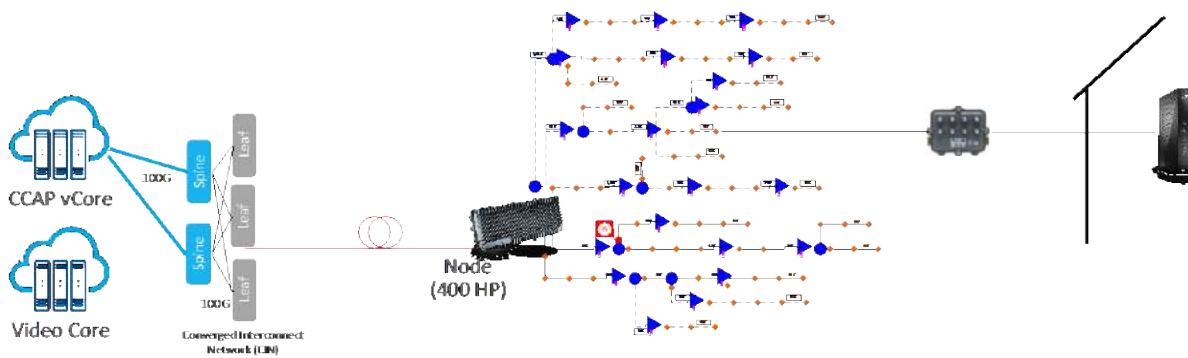
If only 1 to 1.5 Gbps US service tiers are needed, the 1218/204 MHz upgrade is the most cost effective. However, once multi-Gbps US tiers are needed, then the FDX N+0 or the 1.8 GHz N+4 upgrade is required. Our study shows that the FDX N+0 upgrade is roughly twice the cost of the 1.8 GHz ESD N+4 upgrade for this N+4 plant case study.




**Figure 31 – CAPEX estimates for 1218/204 MHz N+4 (both without and with new tap faceplates), 1218 MHz FDX N+0, 1794 MHz ESD N+4 upgrades**



**Figure 32 – DAA N+0 HFC Network Model for 1.2 GHz FDX Upgrade**



**Figure 33 – DAA N+4 HFC Network Model for 1.8 GHz ESD upgrade**



**Figure 34 – Monte-Carlo analysis of “premium” for 1.8 GHz ESD case**

There are other benefits to the ESD approach. For one, it is providing dedicated upstream capacity while FDX is relying on statistical time sharing of US+DS spectrum. The second is that ESD is a much simpler system than FDX which should result in OPEX savings through reduced maintenance and easier to diagnose plant issues.

Because many elements go into an evaluation of this type, a Monte-Carlo analysis was run with various assumptions varying over a reasonable range for the ESD upgrade case. A total of 100,000 Monte-Carlo samples were run to see how much the ESD “premium” might vary. Figure 34 shows that the ESD “expected” premium of 114% expands into a 94% to 134% “confidence interval”.

What is driving the 114% ESD premium, while the FDX N+0 premium is 354%? Figure 31 offers some answers: fiber over lash required for N+0 accounts for 264% of the 354% premium for FDX! The 2<sup>nd</sup> most dominant FDX component is node hardware at ~90%. For ESD, RF amplifiers, labor to replace those, and a material and labor to upgrade tap housings and faceplates to 1.8 GHz account for most of the cost.

For the ESD case, a reasonable “workload” for the whole amplifier and taps upgrade were considered, especially allowing for the whole housing to get replaced. This assumed there is 1 tap per 4 HP; the RF amplifiers are closer to 1 to 20 HP ratio, and nodes are even further – in the assumptions made above it’s 1 to 400 HP. If an amplifier module takes approximately 1 hour to replace and a tap housing takes 45 minutes to replace, how much of an additional labor force would an operator need in order to do these upgrades?

For an MSO serving a region with population of one million people (e.g. ~400,000 HP), Table 4 holds the answer. For the taps portion, it would require an additional full-time crew of ~10, working full-time, over a 5-year period; and doing nothing else. One may describe it as a large task, but 10 additional employees to do this task seems reasonable.

The previous section discussed a corner case where a subscriber with a long RG-6 drop cable at the EoL on a N+5/N+6 stretch plant does not have sufficient network capacity to obtain the highest service tier. Of the half dozen potential options mentioned, the most cost effective would be to replace the drop cable with RG-11. The operator might also consider putting it in a conduit for reliability and ease of later upgrading to a fiber drop as well as consider siamese cable/fiber pair. Adding a mid-span boost amplifier for high frequencies on the last link might triple or quadruple the cost of the drop cable option. And, more

than one mid-span boost amplifiers might be needed on some of the other links in the cascade. Neither of these options pushes fiber deeper so they may not be aligned with the longer strategic direction.

**Table 4 – Estimate of Labor force required to replace amps and taps in 5 years**

	Nodes	RF Amps	Taps	Homes-Passed
Quantity	1,000	20,000	100,000	400,000
Task duration	1 Hour	1 Hour	45 minutes	
24/7 Person-Years	0.1	2.3	8.6	
40hr/week Person-Years	0.5	9.6	36.1	
30hr/week Person-Years	0.6	12.8	48.1	
Required size crew to complete the task in 5 years	<<1	<3	<10	

Trying to jump directly to FTTH can be extremely costly. Since this home is EoL on a N+5/N+6 plant, it might be up to a mile and a half from the fiber node. The costs of pulling fiber can skyrocket, especially if the plant is mostly underground. Some estimates show this to be 100 times the cost of just replacing the drop cable. This is hard to justify for a single customer.

One strategy is to push fiber deeper over a 20-year window to get most customers within 1000' or 1500' of fiber access. At this point, the operator can offer FTTH on demand when and even if needed (e.g. 90%+ of subs may stay on cable 'forever'). To achieve this, the operator can focus on the N+5/N+6 stretch plants in the near-term, with the goal of getting all of their HFC down to N+2/N+3 by the end of the 2020 decade. Then in the decade that follows, the operator can push fiber deeper to achieve N+0/N+1 that can enable FTTH on demand by the end of the 2030 decade. Analysis such as [ULM\_2016] have shown that fiber deep N+0/N+small upgrades are more cost effective than FTTH for HFC brownfield upgrades.

## 8. Conclusion

The title of the paper posed the question “Is “Unity Gain” Still the #1 Objective?” as the cable world migrates to 1.8 GHz DOCSIS 4.0 plant. The paper has shown how much more difficult the 1.8 GHz unity gain task will be compared to 1218 MHz. Looking at stretch plant, the unity gain starts to come up short and the cracks become obvious at the longer cascades like N+5 and N+6.

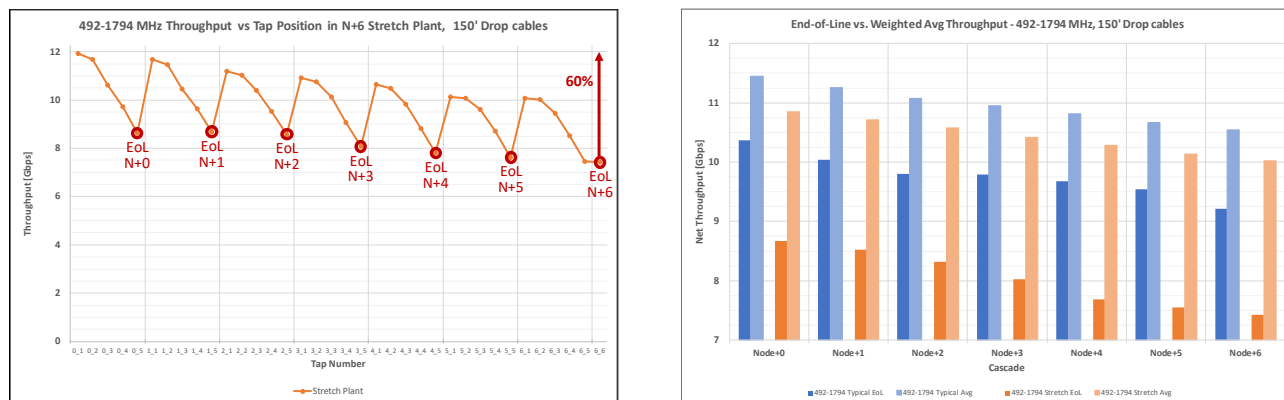
Previously, conventional wisdom would have said that the amplifiers would need to be re-spaced; a potentially expensive proposition. Recently, there has been discussions of adding a mid-span amplifier to boost the gain for these links, which should not be taken lightly, as this adds more active components into the system.

Instead, inherent DOCSIS capabilities can leveraged to compliment unity gain when it comes up short. Previous thoughts of using EoL throughput as an indicator of plant capacity are no longer accurate at 1.8 GHz. DOCSIS 3.1/4.0 OFDM channels with variable bit-loading and multiple profiles enable the CMTS scheduler to maximize the system capacity. This was shown nicely in Figures 18 & 20, which are repeated in Figure 35 below. The N+6 stretch plant saw DOCSIS weighted average capacity gains that were 35% higher than EoL throughput.

The paper explored some of the variables that impact system capacity. In addition to amplifier spacings (e.g. typical vs. stretch plants), it looked at various cable drop lengths and the addition of super-stretched

links into a typical plant. These are all scenarios that allow DOCSIS to optimize the system capacity some more.

The 1.8 GHz ESD plant offers several options for different upstream splits, varying all the way up to 684/834 MHz. Choosing different splits lets the operator balance between upstream and downstream bandwidth. The paper looked at the DS capacities associated with the different splits and then showed the range of DS + US service tier SLA combinations that might be supported near the end of this decade.



The cost and logistic analysis section hopefully gave the reader some sense of the economic tradeoffs that will be encountered when looking at these various options. When the operator finally needs multi-Gbps US tiers, the 1.8 GHz ESD is more cost effective, potentially half of the cost of a FDX N+0 upgrade.

So, to answer the question from our title, the answer is: Maybe Yes! The authors still believe that trying to maintain unity gain is a key objective in any HFC design. It has shown its worth time and time again. However, as it approaches the breaking point, DOCSIS scheduling can fill the gaps on many of these plants on the bubble, reducing the number of HFC links that need drastic action to a significantly smaller amount. The two work well together.

## Acknowledgements

The authors would like to thank Cox Communications and in particular David Job, Principal Eng of OSP Engineering at Cox, for providing the data and insights into the HFC amplifier spacing distribution.

The authors are also extremely grateful to everyone that gave us insights and advice on HFC design and doing a sanity check on what “reasonable” might be. In particular, we want to thank Stuart Eastman of CommScope for sharing some of network design wisdom; and Jay Lazorcik of CommScope for HFC network and unity gain diagrams.

# Bibliography & References

- [ALB\_2019] A. Al-Banna et. al., “Operational Considerations and Configurations for FDX & Soft-FDD,” SCTE Cable-Tec 2019, SCTE
- [CLO\_2019] T. J. Cloonan et. al., “Capacity Planning, Traffic Engineering, And HFC Plant Evolution For The Next 25 Years,” SCTE Cable-Tec 2019, SCTE
- [CLO\_2017] T. J. Cloonan et. al., “The Big Network Changes Coming with 1+ Gbps Service Environments of the Future,” SCTE Cable-Tec 2017, SCTE
- [CLO\_2016] T. J. Cloonan et. al., “Using DOCSIS to Meet the Larger BW Demand of the 2020 Decade and Beyond,” NCTA Spring Technical Forum 2016, NCTA
- [CLO\_2014] T. J. Cloonan et. al., “Simulating the Impact of QoE on Per-Service Group HSD Bandwidth Capacity Requirements,” SCTE Cable-Tec 2014, SCTE
- [CLO\_2013] T. J. Cloonan et. al., “Advanced Quality of Experience Monitoring Techniques for a New Generation of Traffic Types Carried by DOCSIS,” NCTA Spring Technical Forum 2013, NCTA
- [EMM\_2014] “Nielsen’s Law vs. Nielsen TV Viewership for Network Capacity Planning,” Mike Emmendorfer, Tom Cloonan; The NCTA Cable Show Spring Technical Forum, April, 2014
- [FDX\_PHY] “DOCSIS 4.0 Physical Layer Specification”, CM-SP-PHYv4.0-D01-190628, CableLabs 2019
- [FDX\_XSD\_IBC] “Full duplex DOCSIS & Extended Spectrum DOCSIS Hold Hands to Form the 10G Cable Network of the Future”, by F. O’Keeffe et. al., IBC 2019
- [ULM\_2019] J. Ulm, T. J. Cloonan, “The Broadband Network Evolution continues – How do we get to Cable 10G?”, SCTE Cable-Tec 2019, SCTE
- [ULM\_2018] J. Ulm, “Making room for D3.1 & FDX – Leveraging Something Old that is New Again!”, SCTE Journal of Network Operations : Find Fresh Approaches to Plant-Related Topics, Vol 4. No. 1. Dec 2018, SCTE
- [ULM\_2017] J. Ulm, T. J. Cloonan, “Traffic Engineering in a Fiber Deep Gigabit World”, SCTE Cable-Tec 2019, SCTE
- [ULM\_2016] “Adding the Right Amount of Fiber to Your HFC Diet: A Case Study on HFC to FTTx Migration Strategies”, John Ulm, Zoran Maricevic; 2016 SCTE Cable-Tec Expo
- [ULM\_2014] “Is Nielsen Ready to Retire? Latest Developments in Bandwidth Capacity Planning”, John Ulm, T. Cloonan, M. Emmendorfer, J. Finkelstein, JP Fioroni; 2014 SCTE Cable-Tec Expo

# Abbreviations

BW	bandwidth
CAGR	compounded annual growth rate
CAPEX	capital expense
CCAP	Converged Cable Access Platform
CDF	cumulative distribution function
CM	cable modem
CMTS	Cable Modem Termination System
CPE	consumer premises equipment
D3.1	Data Over Cable Service Interface Specification 3.1
D4.0	Data Over Cable Service Interface Specification 4.0
DAA	distributed access architecture
DOCSIS	Data Over Cable Service Interface Specification
DS	downstream
EOL	end of line
EPON	Ethernet Passive Optical Network (aka GE-PON)
ESD	extended spectrum DOCSIS
FDX	full duplex (i.e. DOCSIS)
FTTH	fiber to the home
FTTx	fiber to the 'x' where 'x' can be any of the above
Gbps	gigabit per second
GHz	gigahertz
GOA	grey optics aggregator
GOT	grey optics terminator
HEO	head end optics
HFC	hybrid fiber-coax
HP	homes passed
HW	hardware
I-CCAP	Integrated Converged Cable Access Platform
LDPC	low density parity check (FEC code)
LE	line extender
MAC	media access control
MB	multi-port bridger
Mbps	megabit per second
MHz	megahertz
MSO	multiple system operator
N+0	node+0 actives
NCTA	The Internet & Television Association
OFDM	orthogonal frequency-division multiplexing
OPEX	operating expense
PHY	physical interface
PON	passive optical network
PSD	power spectral density
QAM	quadrature amplitude modulation
QoE	quality of experience



RF	radio frequency
SG	service group
SCTE	Society of Cable Telecommunications Engineers
SNR	signal to noise ratio
STB	set top box
Tavg	average bandwidth per subscriber
TCP	total composite power
Tmax	maximum sustained traffic rate – DOCSIS Service Flow parameter
TX	transmit
US	upstream

# Augmented Intelligence

## Next Level Network and Services Intelligence

A Technical Paper prepared for SCTE•ISBE by

**Dr. Claudio Righetti**

Chief Scientist  
Telecom Argentina S.A.  
Agüero 2392, Buenos Aires, Argentina  
Phone: +5411 5530 4468  
crighetti@teco.com.ar

**Mariela Fiorenzo**

Expert Tech Scientist  
Telecom Argentina S.A.  
mafiorenzo@teco.com.ar

**Omar Hurtado**

Tech Scientist  
Telecom Argentina S.A.  
ojhurtado@teco.com.ar

**Gabriel Carro**

Director Network Technology Architecture and Services  
Telecom Argentina S.A.  
gcarro@teco.com.ar

# Table of Contents

Title	Page Number
Abstract.....	4
Content .....	4
1. Introduction .....	4
1.1. Humans developing AI algorithms.....	4
2. What is Augmented Intelligence and why? .....	5
2.1. Some definitions.....	6
2.2. The goal of Augmented Intelligence .....	7
2.3. Historical Background of Augmented Intelligence.....	8
2.4. Augmented Intelligence use cases.....	9
2.4.1. Healthcare .....	9
2.4.2. Biotechnology .....	9
2.4.3. Financial services .....	10
2.4.4. Retail .....	10
2.4.5. Manufacturing .....	10
2.4.6. Oil and gas .....	10
2.4.7. Geospatial images .....	10
2.4.8. Telecommunications .....	10
3. The technology behind Augmented Intelligence.....	10
3.1. Types of data.....	10
3.1.1. Structured data .....	11
3.1.2. Unstructured data .....	11
3.1.3. Semi-structured data.....	11
3.2. Big Data and Small Data.....	11
3.2.1. Big Data.....	12
3.2.2. Small Data .....	12
3.3. Artificial Intelligence.....	12
3.4. Machine Learning .....	13
3.4.1. Supervised Learning .....	15
3.4.2. Unsupervised Learning .....	15
3.4.3. Reinforcement Learning .....	15
3.4.4. Deep Learning .....	16
3.5. Automation .....	16
3.6. Cognitive computing .....	16
3.7. Intent-Based Networking.....	16
4. Knowledge Plane, state of the art.....	17
4.1. Introduction .....	17
4.2. Autonomic networks .....	18
4.2.1. ETSI (European Telecommunications Standards Institute) .....	18
4.2.2. ONAP (Open Networking Automation Platform).....	20
4.3. Mobile Networks .....	21
4.3.1. O-RAN Alliance (Open Radio Access Network).....	21
4.3.2. 3GPP (3rd Generation Partnership Project).....	21
4.4. Use cases of AI in networks.....	22
4.4.1. General use cases: .....	22
4.4.2. Particular use cases with 3GPP SA2 NWDAF (Network Data Analytics Function).....	24
5. The Future of Augmented Intelligence.....	24
5.1. Augmented Intelligence .....	24
5.2. Explainable AI .....	25
6. Augmented Intelligence at Telecom Argentina.....	25

6.1.	VMAF: a tool for measuring video based on human perception .....	26
6.2.	STEM-ML: a tool for capacity planning .....	26
6.3.	Telecom Argentina Knowledge Plane .....	27
7.	Conclusion .....	28
Abbreviations.....		28

## List of Figures

Title	Page Number
Figure 1 – Gartner Hype Cycle for Emerging Technologies, 2019 (Source: Gartner) .....	6
Figure 2 – Business Value Forecast by AI Type through 2025, expressed in millions of Dollars (Source: Gartner) .....	8
Figure 3 – Three types of data: Structured, Unstructured and Semi-structured Data.....	11
Figure 4 – Artificial Intelligence Timeline: It is not something new.....	13
Figure 5 – Machine Learning tasks divided by the three main categories and most common algorithms. ....	15
Figure 6 – GANA Architecture.....	20
Figure 7 – NWDAF uses case.....	22
Figure 8 – Telecom Argentina Knowledge Plane.....	28

# Abstract

Augmented Intelligence, also known as intelligence amplification, cognitive augmentation, decision support, machine augmented intelligence, and enhanced intelligence, is essentially Artificial Intelligence with a novel approach. While Artificial Intelligence is creating machines to work and react like humans, Augmented Intelligence is using those same machines with a different approach to improve human capabilities. In fact, Augmented Intelligence involves people and machines working together, leveraging their own strengths to achieve greater business value. In other words, the primary goal of Augmented Intelligence is to empower humans to work better and smarter. In this paper, we present the journey we are taking in Telecom Argentina, from data analytics, ML, AI to Augmented Intelligence.

## Content

### 1. Introduction

We started applying data analysis technologies, then machine learning and Artificial Intelligence applied to our networks and services more than a decade ago in the former Cablevisión Argentina and continue after the 2018 merger, in Telecom Argentina. Telecom is in a process of digital transformation and as part of that process we understood that we should not only continue with the application of AI to our mobile, fixed and service networks, but take another step. Focus AI technology on humans.

Many times when people hear about algorithms, robots and AI, they imagine that such technologies compete with them. *“The technology is the easy part. The hard part is figuring out the social and institutional structures around the technology”* [ REF John Seely Brown]

The application of Artificial Intelligence technology may not be successful if it is poorly adapted, designed or implemented. We must ensure that it is designed to help humans think better.

That's why we focus not only on AI technology, but also on human-machine collaboration, processes, and interfaces. This is Augmented Intelligence seeking to elevate human capabilities and experiences. By focusing AgI on serving people rather than replacing them. AgI can help them achieve their greatest potential.

#### 1.1. Humans developing AI algorithms

It is well known that supervised ML algorithms are trained and tested with a large amount of data whose output variables are known (either numerical or categorical). These data consist of a large number of explicative variables or attributes that are in principle chosen at random.

Before "feeding" the machine the data are "worked", eliminating for example the missing data and outliers, standardizing, etc., all this so that no problems in learning the algorithm occur. With the algorithm written by humans, with the data that feeds the training also chosen by humans, we expect that the algorithm can predict a numerical or categorical variable with very little error when an unknown data is presented. The error is inherent in the system because this is an algorithm that does not want to "interpolate" the results but rather "approximate" them and therefore its goal is to minimize the error of its approximation. One way to

fully understand this process is to use the PAC (“probably approximately correct”) Learning framework <sup>1</sup> [ ] .

In a second phase, the human, can diminish the number of explicative variables with which he fed the system and the amount of the same ones, going from a big data to a small data. In short, we are moving towards a process of purification in terms of data quality. The human, expert, could be constantly testing new explicative variables, new sources of data, perhaps unconsciously pursued by the search for a causal law.

But we are not in a world of rules and laws, we are in a world where we are offered an innumerable amount of data with its attributes and we want to use them to decide or to predict at best. But sometimes what "is not" is as important as what "is". What we are not considering because of our own limitations or bias could be being left out of the data that train the algorithm.

And that should be a human task, to go out in search of the unknown, to make them present to us. As we see, human action is immersed in all aspects of AI. Defining the data, building the algorithms, which are constantly evolving and becoming more specialized, interpreting the results.

We could make an analogy with Kant's Copernican revolution, paraphrasing him and say that all knowledge has to "start" with data, but knowledge does not have its only origin in data. The constantly evolving algorithms used do not have their origin in data. Neither does the interpretation we give to the results according to our expertise.

In short, in this phase of our technology, a feedback process is needed between us, the humans and the machines we also create, constantly mediated by the data we know how to get.

As you can see, just like “Artificial Intelligence is not Intelligence, Machine Learning is not Learning” [ref Burkov, Andriy. The Hundred-Page Machine Learning Book (Páginaxvii). Andriy Burkov. Edición 2019].

## 2. What is Augmented Intelligence and why?

Our definition of Augmented Intelligence (AgI) is a powerful intelligence as the result of the collaboration between humans and machines, is the evolution of the Artificial Intelligence (AI) that enables humans to make better-informed decisions from complex data and Machine Learning (ML) algorithms. One aspect of AI that disturbs humanity is the possibility that people may be replaced by machines at their jobs. The goal of augmented intelligence is not to replace human beings or automate them out of existence but to enable them to make better decisions. From our perspective this collaboration between humans and machines will be the enabler for our industry to transition to the information revolution that is coming.

Gartner identifies this emerging technology as key for the design approach of new business solutions, balancing short-term automation with a mid/long-term approach that ensures improving quality not only by automation means, but also by amplifying human talent.

---

<sup>1</sup> In this framework, the learner receives samples and must select a generalization function (called the hypothesis) from a certain class of possible functions. The goal is that, with high probability (the "probably" part), the selected function will have low generalization error (the "approximately correct" part). The learner must be able to learn the concept given any arbitrary approximation ratio, probability of success, or distribution of the samples. The model was later extended to treat noise (misclassified samples). Source: Wikipedia

Figure 1 corresponds to the Hype Cycle for Emerging Technologies that Gartner publish every year, since 1995. According to the cycle, we can see that this technology will reach the plateau in 2 to 5 years [ref: [gartner.com/SmarterWithGartner](http://gartner.com/SmarterWithGartner)].



**Figure 1 – Gartner Hype Cycle for Emerging Technologies, 2019 (Source: Gartner)**

## 2.1. Some definitions

According to Merriam-Webster, information can be defined as the “communication or reception of knowledge or intelligence; knowledge obtained from investigation, study, or instruction”. The future information revolution will be conducted by three main pillars: Free data and information, Small data and Augmented Intelligence (ref en Bell Labs). In this future, the goal will be to store the right amount of data that enables data scientist to discover knowledge and extract useful information by using new tools provided by the AgI.

We found that there is a confusion about the meaning of the terms data and information, it is believed that they both mean the same. These are the definitions we considered for this paper:

- *Data*: the actual captured observations.
- *Information*: the determination of relationships between data.
- *Knowledge*: the determination of models that describe the meaning of Information.
- *Technology*: a manner of accomplishing a task especially using technical processes, methods or knowledge (ref: Webster’s Dictionary 2015).

Although data is easy to obtain, to establish the relationships between data is not. In many cases, it is possible to determine this relationship as the statistical correlation but it not necessary implies causation. Determining models and applications where this correlation makes sense is the most important component to succeed in discovering the meaning of information, that is, knowledge. We can define the five basic

knowledge acquisition questions: *who, what, when, where* and *why*? These questions are an ideal starting point for assessing the relative value of information.

But the question *why*? is distinct from *who, what, when* and *where*? because forming an answer to *why*? implies some understanding or knowledge of a situation.

## 2.2. The goal of Augmented Intelligence

How this human-machine collaboration works? Both humans and machines have limitations doing their tasks. Humans cannot process or understand huge amount of data or complex data such as metadata, images, videos, etc. that machines can do and in a very fast way. On the other hand, machines cannot make business or economical decisions or understand the quality of data that they are processing or to choose the best strategy to achieve the industry goals.

Therefore, this collaboration surpasses humans and machines limitations by combining their intelligence and strengths. Humans are the ones that must use their knowledge and expertise to redesign the business process because the technology is not going to make such difficult decisions to transform the industry.

Augmented Intelligence will produce outcomes that neither humans nor machines could achieve alone (IBM). There are some steps to implement AgI:

- Decide whether to change the business process and task flow for human-machine cooperation<sup>2</sup>.
- Select which tasks and decisions within the business process to automate.
- Determine the proper AI tools.
- Determine what data to acquire to better understand and model the business and customers.
- Build the data models.
- Test the results for reliability and accuracy.

Businesses powered by AgI systems have these four key attributes in common [ref: AgI eBook]:


- Discover and learn from hidden meaning in data and user interactions.
- Continuously learn, evolve, and improve with time.
- Build and test new user and process engagement capabilities.
- Drive new business processes and business model innovation.

In the last years, AgI was ranked in second place in AI technology rankings in terms of the value they create for businesses. However, Gartner predicts that “Decision support and AgI will surpass all other types of AI initiatives”, as it is shown in Figure 2, doubling virtual agents by 2025.

---

<sup>2</sup> We will see later in section 2.3 the importance of this point in the power of AgI





**Figure 2 – Business Value Forecast by AI Type through 2025, expressed in millions of Dollars (Source: Gartner).**

### 2.3. Historical Background of Augmented Intelligence

Throughout the history of humanity, man has created tools (language, writing elements, mathematics, etc.) that have allowed him to increase his intelligence. In the early 1960s Doug Engelbart<sup>3</sup> was researching how those tools shape our thoughts, founding the field of human–computer interaction. At the time, most of his colleagues only viewed computers to compute numbers and somehow man machine competition. However, he saw something deeper: He saw a way to increase the human mind [ ref Engelbart, D. C. (1962, October). “Augmenting Human Intellect: A Conceptual Framework.” Retrieved 10 July 2020 from [https://www.dougenelbart.org/pubs/papers/scanned/Doug\\_EngelbartAugmentingHumanIntellect.pdf](https://www.dougenelbart.org/pubs/papers/scanned/Doug_EngelbartAugmentingHumanIntellect.pdf)]

On February 10, 1996, a computer won a game of chess against a world champion for the first time. The computer was Deep Blue, a machine designed by IBM. An improved version of Deep Blue recorded its famous May 11, 1997 victory over world champion Kasparov, a milestone in artificial intelligence. Designed to understand high-power parallel processing, the "brute force" system could examine 200 million chess positions per second, beating Grandmaster 3.5-2.5. The story quickly centered around a Man vs. Machine narrative. For Kasparov it was a turning point. He may have asked himself: "Why want to compete against a machine when we could play with a machine?"

The next year, Garry Kasparov held the world’s first game of “Centaur<sup>4</sup> Chess” [ref <https://www.parc.com/blog/half-human-half-computer-meet-the-modern-centaur/>] or, as it is more commonly known today, “freestyle chess”<sup>5</sup> (The concept of using computers to augment play had been around for a long time.) Humans can use input from chess programs to select their moves.

<sup>3</sup> Was an American engineer and inventor, and an early computer and Internet pioneer. He is best known for his work on founding the field of human–computer interaction, particularly while at his Augmentation Research Center Lab in SRI International, which resulted in creation of the computer mouse, and the development of hypertext, networked computers, and precursors to graphical user interfaces. These were demonstrated at The Mother of All Demos in 1968. Engelbart's law, the observation that the intrinsic rate of human performance is exponential, is named after him” <https://www.dougenelbart.org/>

<sup>4</sup> A centaur (/ˈsɛntɔːr/; Greek: κένταυρος, kentauros, Latin: centaurus), or occasionally hippocentaur, is a creature from Greek mythology with the upper body of a human and the lower body and legs of a horse.

<sup>5</sup> In Freestyle Chess, human players are assisted by computers, software, and database tools.

In 2005 a freestyle chess tournament was organized, team called ZackS won by beating an opponent that included Vladimir Dobrov, a grandmaster, his highly rated teammate, and their computer programs.

The two members of the team were amateur chess players. They didn't play with the best hardware in the world. In fact, they had three different AI systems that run on mass-use computers. We had average players, with average computers, but a very good workflow. One of them was a soccer coach, the other was a database administrator. In this team the important thing was how they interacted and collaborated with the machines. That interface is what allowed it to succeed.

“Weak human + machine + superior process was greater than a strong computer and, remarkably, greater than a strong human + machine with an inferior process.” [ref Garry Kasparov, *How Life Imitates Chess: Making the Right Moves—from the Board to the Boardroom* (New York: Bloomsbury, 2007), pp166]

By combining human intelligence with technological intelligence (Augmented Intelligence), these players tend to outdo anyone. In other words, centaurs can outperform humans and machines in the chess domain.

Summarizing centaurs combine the main characteristics of humans: INTUITION, JUDGMENT and FLEXIBILITY. With those of the machines: CONSISTENCY, PRECISION, SCALABILITY. Machines are for giving ANSWERS and humans are for asking QUESTIONS.

### *Centaur > Man or Machine*

The “platforms”<sup>6</sup> have long since incorporated centaurs, today we say Augmented Intelligence, into their work teams.

## **2.4. Augmented Intelligence use cases**

Although it is early days, AgI has already had a positive impact on many sectors and the nature of this technology, which learns more, adapts faster and continuously improves with time, means that early adopters do gain an advantage. But the main source of augmentation's business value will continue to be an improved human experience. AgI will deliver a level of personalization and it will also minimize errors, creating a higher standard of service.

Some of the following examples correspond to real applications and others are just possible use cases.

### **2.4.1. Healthcare**

Augmented intelligence is transforming the industry, from detecting outbreaks to providing more customized care and explainable diagnoses. It has also been used in the fight against COVID-19.

### **2.4.2. Biotechnology**

Biotech companies help augment doctors for radiology and clinical data. They use machine learning algorithms to extract key characteristics from radiology and pathology images and also organize the unstructured data of each patient's clinical notes, test results and health history. Doctors are provided with information to combine with their expertise in selecting treatment options for the patient.

---

<sup>6</sup> The Business of Platforms: Strategy in the Age of Digital Competition, Innovation, and Power

### **2.4.3. Financial services**

Assisting financial planners to offer personalized services based on the customer's goals, capacity and risk appetite.

### **2.4.4. Retail**

AgI helps increasing shopper engagement and conversion by enabling online shoppers to shop the way they think, using machine cognition of their declared, observed, and inferred behaviors.

### **2.4.5. Manufacturing**

Aiding and accelerating the generative design process, whereby a human worker inputs the parameters and the machine finds countless ways of designing the object. The machine explores a plethora of options in record time and the human uses their expertise to select the best option, delivering value to customers and boosting efficiency.

### **2.4.6. Oil and gas**

Optimizing precision drilling. The human worker can understand the environment they are operating in more accurately, leading to faster results and less wear, tear and damage to machinery. The list of possibilities is endless, but the common element is clearly to increase efficiency by heightening the worker's knowledge.

### **2.4.7. Geospatial images**

Machine Learning is used to analyze geospatial imagery data. This provides real estate, energy and government agencies with information on land use, car and air traffic, and demographic trends so they can make better decisions.

### **2.4.8. Telecommunications**

AgI uses information gathered from applications to maximize network configurations and simplify troubleshooting. This next level of network intelligence comes from AI, data analytics and ML, to enable better correlation among events on the network, user and device behavior.

## **3. The technology behind Augmented Intelligence**

There is a convergence of technologies that have come together to lead the market towards the reality of augmented intelligence and in this section, we will explore them.

### **3.1. Types of data**

Augmented Intelligence systems can work with all types of data (structured, unstructured, semi-structured and metadata) from many sources, across disparate and siloed systems.



**Figure 3 – Three types of data: Structured, Unstructured and Semi-structured Data**

### **3.1.1. Structured data**

It is considered the most ‘traditional’ form of data storage. It refers to information that has a defined length and format such as numbers, names or dates. This type of data is linked to a pre-defined data model and is therefore straightforward to analyze. They fit to a tabular format with relationship between the different rows and columns. This structure made it possible to create understandable answers to questions inside this data [ref: <https://www.bigdataframework.org/data-types-structured-vs-unstructured-data/>]

### **3.1.2. Unstructured data**

It does not have a predefined data model nor is it organized in a predefined manner. Therefore, it is not surprising that most of the information in the world is unstructured, for example, videos, images, text documents. The ability to analyze unstructured data is especially relevant in the context of Big Data, since a large part of data in organizations is unstructured. Of course, there is inherent structure, but the difference is that humans must do the hard work to understand the hidden structure of the data.

### **3.1.3. Semi-structured data**

It is a form of structured data that does not conform with the formal structure of data models associated with relational databases or other forms of data tables, but nonetheless contain tags or other markers to separate semantic elements and enforce hierarchies of records and fields within the data. Therefore, it is also known as self-describing structure. JSON and XML are two examples of this type of data. They are considerably easier to analyze than unstructured data and many tools have the ability to ‘read’ and process them.

## **3.2. Big Data and Small Data**

The terms Big Data and Small Data have become popular buzzwords over the last years and it is not always clear what either of these terms means or how or when to use each.

The acquisition of information requires the ability to capture data, compute something based on the data available and obtain a result.

Over the last decade, the promising concept of Big Data has generated huge expectations to industries. Therefore, companies have purchased and deployed scalable storage and processing systems with the intention of preserve every single byte of data obtained from their systems and customers. But actually, much of this data has no real information or valuable content. As storage costs began to increment, big data applications suffered a transformation into a new type of applications of small data where the value arises not from the volume of the data set, but from the ability to extract useful information and to make decisions

based upon the smallest data set. The goal will not be to measure and store every byte; it will be to measure and store “just the right amount” of data [ref Bell Labs].

Next, we present some formal definitions with respect these concepts.

### **3.2.1. Big Data**

Gartner proposed an early definition using three ‘Vs’ (Volume, Velocity, and Variety) to represent the key characteristics of Big Data. Certainly, it is still widely used.

“Big Data is high-volume, high-velocity and/or high-variety information assets that demand cost-effective, innovative forms of information processing that enables enhanced insight, decision making, and process automation.”

More recent definitions include other “V-terms” into the Gartner model, adding Veracity to reflect data accuracy and Value to address the usefulness of the data.

Big data techniques are designed to manage a huge volume of disparate data at the right speed and within the right time frame. The goal is to enable near real-time analysis and action.

### **3.2.2. Small Data**

A formal definition of small data has been proposed by Allen Bonde, former vice-president of Innovation at Actuatet:

“Small data connects people with timely, meaningful insights (derived from big data and/or “local” sources), organized and packaged – often visually – to be accessible, understandable, and actionable for everyday tasks”.

Martin Lindstrom defines it as ‘the tiny clues that uncover huge trends’, based on observational data. It’s also defined as data that is small enough size for human comprehension.

From the white paper [ref Philosophy of Small Data], “Just only one apple falls on Isaac Newton’s head, not ten, not thousand”.

## **3.3. Artificial Intelligence**

Artificial Intelligence means getting a computer to mimic human behavior in some way. The goal is to get computers to perform tasks as human: things that required intelligence.

Despite recent hype around the technology, AI is not a new technology and not a product of this century’s innovations. The beginnings of AI can be traced to the middle of the 20th century. During the second world war, an English mathematician and computer scientist, Alan Turing documented his ideas on creating an intelligent machine. He proposed a test for machine intelligence:

A human evaluator would judge natural language conversations between a human and a machine designed to generate human-like responses. The evaluator would be aware that one of the two partners in conversation is a machine, and all participants would be separated from one another. The conversation would be limited to a text-only channel such as a computer keyboard and screen so the result would not depend on the machine's ability to render words as speech. If the evaluator cannot reliably tell the machine from the human, the machine is said to have passed the test. The test results do not depend on the machine's

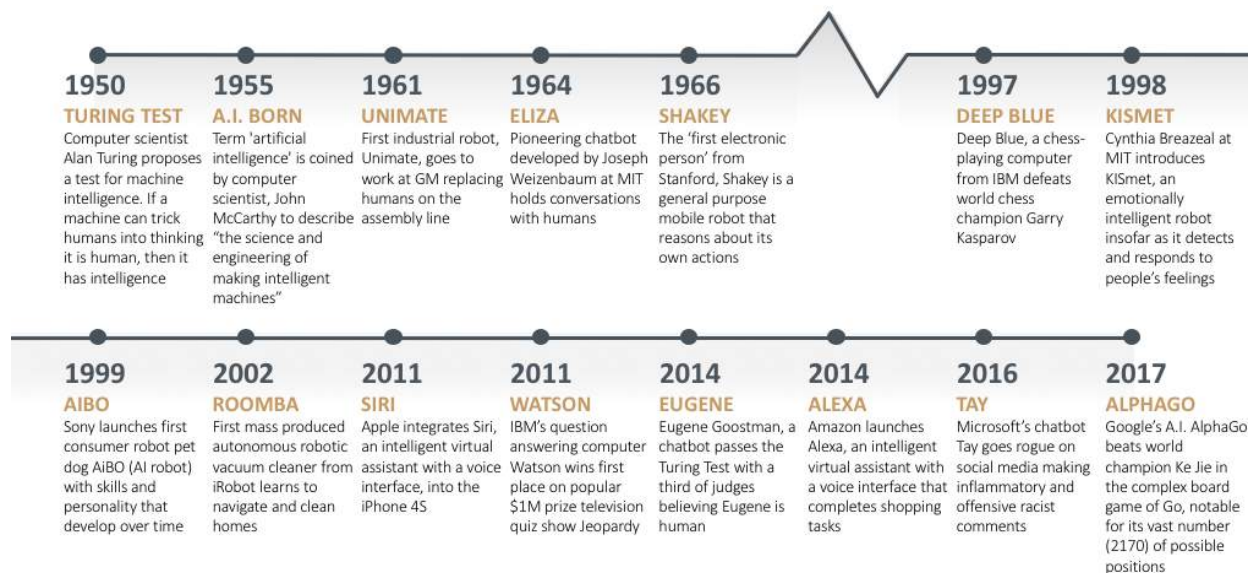
ability to give correct answers to questions, only how closely its answers resemble those a human would give [ref Wikipedia Turing Test].

In short, if a machine can trick humans into thinking it is human, so then it has intelligence.

Marvin Minsky of MIT's Project Mac, who has made major contributions to Artificial Intelligence, in 1970 said: "In from three to eight years we will have a machine with the general intelligence of an average human being. I mean a machine that will be able to read Shakespeare, grease a car, play office politics, tell a joke, have a fight. At that point the machine will begin to educate itself with fantastic speed. In a few months it will be at genius level and a few months after that its powers will be incalculable."

Figure 4 illustrates the AI Timeline. Despite the fall between 1966-1997, which was called the AI winter, the constant evolution of this technology is evident.

Today, AI refers to a range of technologies from automation to deep learning. It includes the subfields of natural language processing, vision, robotics, machine learning, and knowledge representation and reasoning.



**Figure 4 – Artificial Intelligence Timeline: It is not something new.**

The question that arises is what the difference between Artificial Intelligence and Augmented Intelligence is. The answer can be found within its objectives; the goal of an artificial intelligence system is to simulate human cognitive capabilities in a system that can function independently of humans. In contrast, the goal of an augmented intelligence system is to enhance human intelligence by human-machine collaboration to get work done.

### 3.4. Machine Learning

Machine learning is a form of artificial intelligence that enables a system to learn from data rather than through explicit programming of a set of rules. It consists of a variety of types of algorithms, all of which learn from data. A machine learning model is the output generated when you train your machine learning algorithm with this data. After training, when you provide the model with new data input, its output will be providing new insights such as data classification or prediction.

The end-to-end machine learning process includes the following phases:

- Business Goal Identification
- ML Problem Framing
- Data Collection and Integration
- Data Preparation
- Data Visualization and Analytics
- Feature Engineering
- Model Training
- Model Evaluation
- Business Evaluation
- Production Deployment

Machine learning models can be online or offline, online models are constantly ingesting data and interacting with it in near real time mode improving the model outcomes. On the other hand, offline models once they are deployed, they only can be retrained manually with new data.

There are several approaches to machine learning that are relevant to the ability to create algorithms that support the industry problems, they are based on the type of the data. These approaches are divided into three main areas: supervised learning, unsupervised learning (it includes deep learning) and reinforcement learning.

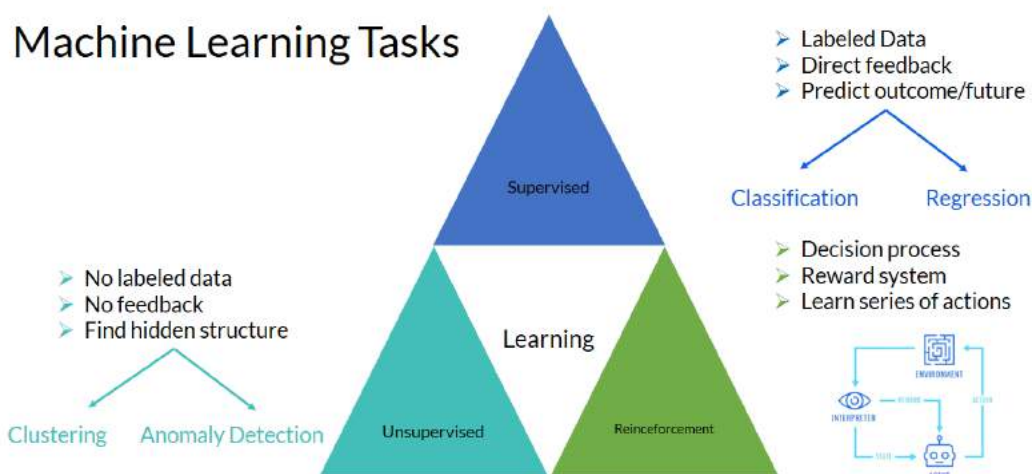
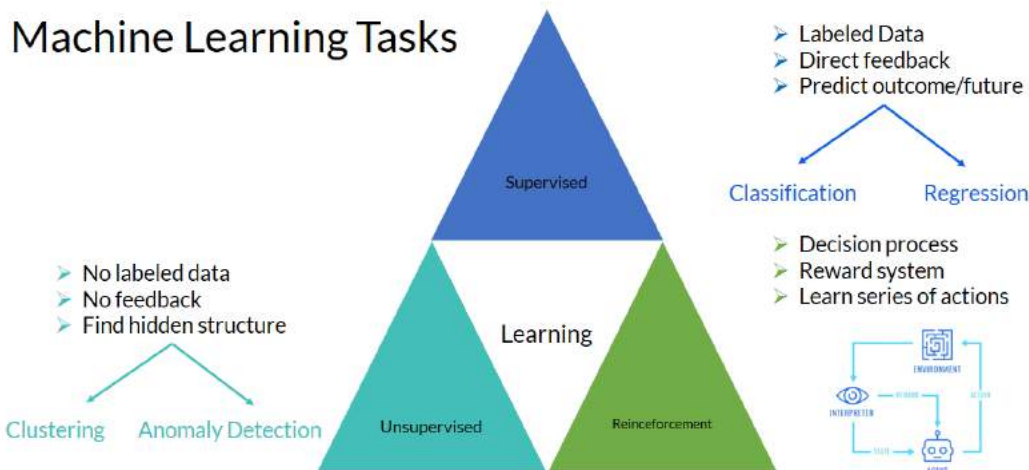


Figure 5 gives an explanation about the type of data, type of outcome and the algorithms methodology of each category. It also shows some of the most common techniques: regression, classification, anomaly detection, etc.





**Figure 5 – Machine Learning tasks divided by the three main categories and most common algorithms.**

### **3.4.1. Supervised Learning**

Supervised machine learning algorithms are designed to learn by example. During training, the algorithm will search for patterns in the data that correlate with the desired outputs. After training, a supervised learning algorithm will take in new unseen inputs and will determine which label the new inputs will be classified as based on prior training data. The resulting model must be evaluated against test data to see how well it learned. If the model is fit to only represent the patterns that exist in the training set, overfitting<sup>7</sup> occurs. Using unseen data for the test set can help you evaluate the accuracy of the model in predicting outcomes.

### **3.4.2. Unsupervised Learning**

Unsupervised learning is a set of statistical tools for scenarios in which there is only a set of features and no labels. With these techniques we are interested in finding discovering subgroups of similar observations. It tends to be more challenging, because there is no clear objective for the analysis. Besides, it is hard to evaluate if the obtained results are good, since there is no accepted mechanism for validating results on an independent dataset, because we do not know the true answer.

### **3.4.3. Reinforcement Learning**

Reinforcement learning is a behavioral learning model. It is about taking suitable actions to maximize reward in a situation wherein an agent interacts with a new environment using actions and discovering errors or rewards. It is employed by various software and machines to find the best possible behavior or path it should take in a specific situation. In short, a reinforcement learning system learns through trial and error. Therefore, a sequence of successful decisions will result in the process being “reinforced” because it best solves the problem at hand. Gaming and robotics are the most common applications of this technique.

<sup>7</sup> The production of an analysis that corresponds too closely or exactly to a particular set of data, and may therefore fail to fit additional data or predict future observations reliably [ref: Oxford Dictionaries.com]



#### **3.4.4. Deep Learning**

Deep learning is a specific method of machine learning that incorporates neural networks in successive layers in order to learn from data in an iterative manner. Deep learning is especially useful when you are trying to learn patterns from unstructured data.

A neural network consists of three or more layers: an input layer (ingested data), one or many hidden layers (which include weighted nodes), and an output layer (the outcome). The term deep learning is used when there are multiple hidden layers within a neural network.

Neural networks and deep learning are often used in applications where images, videos or speech are involved.

### **3.5. Automation**

Within MSOs and MNOs, this means automation of processes that were previously carried out by people like configuration, management, operation and testing of physical and virtual devices within the network. With growing costs and the daily emergence of bandwidth-hungry applications, networks cannot be managed manually. Increased levels of network automation help to reduce complexity and are essential for businesses to keep up in the digital world. AI is an enabling technology that may (or may not) help with the process of automation. What it culminates is a network that is highly predictable and highly available improving the business outcomes.

### **3.6. Cognitive computing**

Like AI, cognitive computing is based on the ability of machines to sense, reason, act and adapt based on learned experience. Cognitive computing refers to computing that focuses on reasoning and understanding at a higher level and in a manner that is analogous to human cognition, rationale, and judgement. Applications of cognitive computing include speech recognition, sentiment analysis, face detection, risk assessment and fraud detection. The difference between AI and cognitive computing lies in the way they approach the purpose of simplifying tasks. AI is used to augment human thinking and solve complex problems. Cognitive computing mimics human behavior and reasoning to solve complex problems similar to the way humans solve problems.

### **3.7. Intent-Based Networking**

“Intent” is the keyword in this technology, which describes a network’s business objective or an outcome.

Intent-based networking (IBN) is an emerging technology concept that aims to apply a deeper level of intelligence and intended state insights to networking. Ideally, these insights replace the manual processes of configuring networks and reacting to network issues. The goal is networking that uses machine learning and cognitive computing to enable more automation and less time spent on manual configuration and management. With intent-based networking, network administrators define the intent and the network’s software finds how to achieve that goal using AI and ML by performing routine tasks, setting policies, responding to system events, and verifying that actions have been done.

These systems not only automate time-consuming tasks and provide real-time visibility into a network’s activity to validate a given intent, they also predict potential deviations to that intent,

and prescribe the action required to ensure it. This greater intelligence makes the network faster and more agile and reduces errors [ref <https://www.vmware.com/topics/glossary/content/intent-based-networking#:~:text=Intent%2Dbased%20networking%20relies%20on,and%20actions%20have%20been%20achieved.>].

## 4. Knowledge Plane, state of the art

The research community has considered in the past the application of AI techniques to control and operate networks. For example, in 2003 David Clark et. al propose the knowledge plane (KP) as a *pervasive system within the network that builds and maintains high level models of what the network is supposed to do, in order to provide services and advice to other elements of the network. The knowledge plane is novel in its reliance on the tools of AI and cognitive systems* (Clark, Partridge, Ramming, & Wroclawski, 2003).

The Knowledge Plane (KP) paradigm proposes the evolution to a cognitive network, where the devices learn, decide, and act to achieve end-to-end goals. This emerging paradigm is clarifying a set of new cognitive-based protocols and algorithms that optimize network's performance.

In (Mestres, y otros, 2017) progress is made in the definition of a new paradigm based on this plane. This is knowledge-defined network (KDN) operates by means of a control loop to provide automation, recommendation, optimization, validation and estimation.

### 4.1. Introduction

There are several organizations and working groups proposing frameworks, standards and how to apply the AI to the industry of communications, networks and services. Below we detail those in which Telecom Argentina is participating or is in consultation with:

- SCTE & CableLabs AI/ML Working Groups
- Telecom Infra Project - AI/ML Working Group
- TM Forum - AI & Data Analytics
- ITU-T Study Group 13 and ITU-T FG ML5G Studying network architectures, use cases, and data formats for the adoption of machine learning methods in 5G and future networks.
- ETSI ISG ENI (Experiential Network Intelligence). Defining a cognitive network management architecture based on AI methods and context-aware policies; five deliverables have already been released
- 3GPPP
- ONAP

AI began to be studied in most of the different Standards Developing Organizations (SDO). We provide in this item a summary of what they are proposing, how they are defining it and the specific use of this new technology in each branch of the networks.

As has been said repeatedly in various academic circles, AI makes use of another phenomenon that occurs: the massive growth of data, driven mainly by network technologies such as IoT, 5G and 10G project.

The introduction of AI and automation using AI, in short, seeks to ensure greater performance and efficiency of networks.

The paradigm shift that will be brought about by the introduction of these new technologies includes a substantial shift from a focus on network operations to a focus on the user experience.

We have conceptualized these AI and automation tools that interact with the different types of networks as **Knowledge Plane**, a "place" where the massive amount of data obtained from the network is processed with the different AI tools, either in real time or in a post-processing, and that, based on results, produces modifications in the network itself. This is called **closed-loop automation**.

Before we start it would be necessary to break down what we used to call "network" into a general form.

#### Mobile networks:

Split up into:

- Terminals (mobile phones, IoT device, etc.)
- Access (cells, RAT, Fronthaul/Midhaul/Backhaul)
- Packet core (EPC (4G), 5GCN, etc)

#### Fixed networks:

- XDSL (DSLAM/BNG)
- FTTH, GPON (ONT/OLT)
- HFC (CM/CMTS)
- IPBB (routers, DNS, CDN, etc.)

And in common with all of them, a growing trend towards virtualization.

## **4.2. Autonomic networks**

### **4.2.1. ETSI (European Telecommunications Standards Institute)**

In the current research on the ETSI recommendation we still do not see clearly the difference between the initiatives that are established within this SDO, as some of them promote doing exactly the same thing. This is further complicated by the fact that each of them does not refer to interaction with the others. Anyway, we consider it important to explain them.

#### **4.2.1.1. ETSI-ENI (Experiential Networked Intelligence)**

The concept of ENI refers to a working group developed at ETSI to improve the operator's experience that is active since 2017.

This working group has collaboration with all the SDO's in the industry.

It is based on the introduction of AI systems in the Network Management System and a control model that takes the following actions:

- Observe
- Guide
- Decide
- Act

The **knowledge plane**, in ETSI is the ENI, which is materialized in this recommendation as a layer that covers transversally all the layers of the MANO (NFV Management and Orchestration) architecture.

In this case the AI-based system will have to adjust the services offered according to the following possible changes:

- In the client's needs
- Under the ambient conditions
- In business objectives

As a challenge it is established:

- Determine which services comply with the SLA and which would be about to fail to comply depending on the changing context.
- Provide an experiential architecture (i.e. an architecture that uses the benefits of AI).
- Establish incident detection.
- Possessing the capacity for autonomous incident management.

#### **4.2.1.2. ETSI-ZSM (Zero-touch service and network management)**

In 2017 ETSI launched the Zero touch network and Service Management Industry Specification Group (ZSM ISG).

The aim of this group is to specify all processes and operational tasks as an example:

- Deployment
- Configuration
- Assurance
- Optimization

are executed automatically (agile, efficient management and automation).

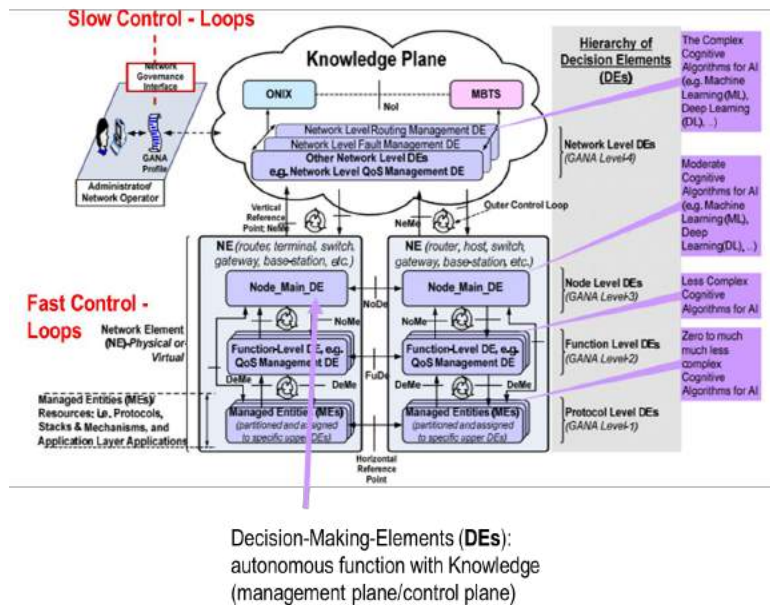
#### **4.2.1.3. ETSI-GANA (Generic Autonomic Networking Architecture)**

It is a reference architecture model for the creation of autonomous networks, "cognitive" networks and self-management of networks and services in which the AI plays an important role in such autonomy.

It is designed to realize the AMC (Autonomic Management and Control) paradigm (closed-loop service instantiation and adaptive operations) of networks and services with reference to:

- Control architectures
- Management architectures

The relevance of each of these ETSI proposals needs to be analyzed to see if they replace or complement each other.



**Figure 6 – GANA Architecture**

#### 4.2.2. ONAP (Open Networking Automation Platform)

ONAP provides a comprehensive platform for the real-time, policy-based orchestration and automation of physical (PNF: Physical Network Functions) and virtual (VNF: Virtual Network Functions) network functions that will enable software, network, IT and cloud providers and developers to rapidly automate new services and support full life-cycle management.

In short, ONAP is a platform that offers a complete set of tools to automate assurance processes in the field of network management.

Some operators have selected it to automate layers of:

- The MANO domain (except for the VIM functions (VNF life cycle management)).
- The orchestration service for both PNFs and VNFs.

ONAP makes use of a functionality called DCAE (Data Collection, Analytics and Events Project) that is the general name for several components that collectively fulfill the role of Data Collection, Analysis and Event Generation.

The DCAE architecture aims at the deployment of components and the composition of flexible, pluggable, microservice-oriented and model-based services. DCAE also supports multi-site data collection and analysis that are essential for large ONAP deployments.

The DCAE is a place where analytics applications and AI/ML models could reside.

### **4.3. Mobile Networks**

#### **4.3.1. O-RAN Alliance (Open Radio Access Network)**

This SDO, which is "openness", proposes exactly the same solutions as the SDOs in terms of the use of the AI, but specifically applied to the RAN.

It proposes that networks should:

- Be self-managed.
- Be able to take advantage of new learning-based technologies to automate the operational functions of the network and reduce OPEX.
- Leverage emerging deep learning techniques to integrate intelligence into every layer of the RAN architecture.

Embedded intelligence, applied at both the component (NE) and network levels, will enable dynamic local allocation of radio resources and will optimize the efficiency of the entire network in closed-loop automation using AI.

As it is possible to see, the O-RAN describes in the orchestration and automation layer a Non-Real Time RAN Intelligence Controller (RIC) with AI/ML capabilities.

It should be remembered that there are currently initiatives for the virtualization of the Non-Real Time BBU.

#### **4.3.2. 3GPP (3rd Generation Partnership Project)**

Service and System Aspects (SA) is the Technical Specifications Group of 3GPP where most work is currently being done on the use of AI, especially in

- SA WG2, Architecture
- SA WG5, Telecom Management

##### **4.3.2.1. SA2 NWDAF (Network Data Analytics Function)**

The NWDAF was first introduced into the 5G system architecture at the 3GPP SA2#119 meeting in February 2017.

The NWDAF (Network Data Analytics Function), as defined in TS 23.503, is used to collect data such as FCAPS (Fault Configuration Accounting Performance Security) events and Data events, and then perform the analysis of the data centrally. An NWDAF can be used for analysis of one or more network slices.

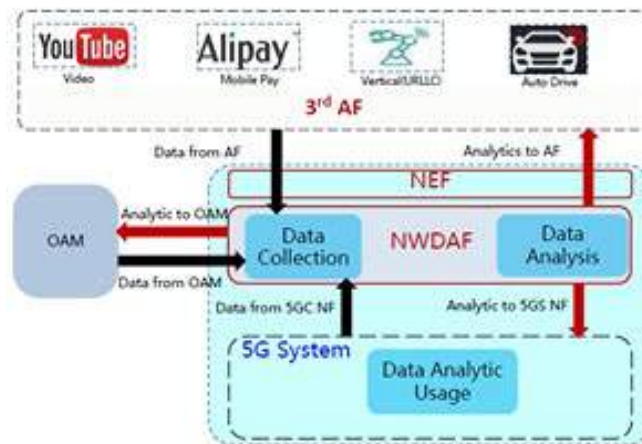


Figure 7 – NWDAF uses case

#### 4.3.2.2. SA5 MDAF (Management data analytics function)

MDAF (management data analysis function) uses network management data collected from the same network, for example data related to

- Services
- Slicing functions
- Network functions
- Related slicing and networking functions

and perform the corresponding "analytics".

The MDAF can be deployed at different levels, for example, at the domain level:

- RAN Management Data Analytics Function
- Network Management Data Analytics function
- Network Slice Subnet Instance Management Data Analytics function

In order to establish a difference between NWDAF and MDAF we can say that the first one is a function in the Packet Core domain or for the management of use cases in the slice domain, while the last one is in the OAM layer supporting the assurance. Both relate to each other by exchanging information.

### 4.4. Use cases of AI in networks

#### 4.4.1. General use cases:

The following cases are added to those already established in each standard:

- In the NOC (Network Operation Center):
  - In the preventive support to allow to identify and solve problems before they affect the performance of the network, avoiding critical cuts and providing stability.

- Incident pattern recognition.
  - Incident based subscriber type clustering.
- Anomaly detection
  - Discovery of deviations from standard behavior.
  - Determination of outliers is multidimensional spaces.
  - AI Tools:
    - ARIMA (autoregressive integrated moving average)
    - Random Forest
    - Self-encoder
    - Principal Component Analysis
- In the RAN
  - Embedded Analytics components (incorporation of chips) in the same Radio Base Stations to perform:
    - Closed-loop automation (internal automation and localized training components).
    - Improve performance and spectral efficiency.
    - Improving Mobility Management.
    - Adaptation of links.
    - Energy saving in MIMO (Multiple Input Multiple Output) in its sleep state.
    - Reduce energy consumption.
    - Evolution of SON (Self-organizing network).
    - Interference diagnosis.
    - Real-time analysis in the baseband unit.
    - ML to improve the algorithms of the cell itself:
      - In the User Plane.
      - With QoE optimization.
      - In the management of radio resources.
      - In the scheduler (especially for MIMO).
- On terminals connected to the mobile network.
- IoT terminals.
- AI can improve network efficiency by being able to cluster and detect anomalies in the initial state of implementation of the ever-growing diversity of devices, in an environment where vendors freely interpret GSMA standards.
- Cellphones.
- The application of the AI tools as in the case of the RAN network elements, will allow them to
  - Reduce latency time and optimize spectrum management.
  - Applications are developed where they can be made:
    - Image and sound recognition and interpretation.
    - Development of Augmented Reality and Virtual Reality.
- Visual inspection of network equipment (deep learning).
- In Energy.
- In Planning:
  - Work is already underway with the mobile access management of Network and Service Planning using forecasting methods (e.g. ARIMA).



- On the WiFi:
  - Not yet determined by any organism.

#### **4.4.2. Particular use cases with 3GPP SA2 NWDAF (Network Data Analytics Function)**

The NWDAF can be used for the following applications:

- Assist in the provision of Quality of Service (QoS) profiles.
- Assist in the adjustment of the quality of service (QoS).
- Assist in policy determination.
- Collaborate with 5G Edge Computing.
- Improve performance and monitoring of mIoT (massive Internet of Things) terminals.
- Assist in load/balance balancing of NF (Network Function).
- Assist in areas of the network with instability (oscillating conditions).
- Improve performance and monitoring of mIoT terminals.
- Support for the exposure of the network status on the Northbound interface using APIs (application programming interface).

## **5. The Future of Augmented Intelligence**

### **5.1. Augmented Intelligence**

As we mention before Augmented Intelligence is still in the technology trigger phase in the hype cycle according to Gartner. And, in early stage of the evolution of Artificial Intelligence and Machine Learning. The main question that organizations is asking about the future of these technologies is: Will the human-machines collaboration result in fewer jobs for people?

What is clear is that the nature of work is already changing and will continue to change through the human-machine collaboration. Both, humans and machines will do what they do best. Machines will automate routine tasks that don't need human intelligence to let humans focus on handling exceptions. When humans handle exceptions, they must get an informative alert with context from the machine, often with a recommendation on how to proceed.

There will be many new jobs that do not exist today. With the evolution of augmented intelligence and its presence in more and different domains, there will be a greater need for regulatory frameworks. In the future, many jobs will be needed to manage augmented intelligence and handle the exceptions have not existed before. One of the greatest challenges for society will be the massive training of those who are displaced by intelligent systems to fill the new jobs that augmented intelligence enables.

Also, in the future, data and information will become free and freely available, big data will be replaced by small data to discover “knowledge”, and new augmented intelligence tools will be developed that assist in the acquisition of knowledge (cognition) by enabling critical thinking from multiple perspectives.

## 5.2. Explainable AI

Another emerging field is Explainable AI (XAI). Whether when an expert must decide based on augmented intelligence tools or when they are implemented in our networks and services, they must understand how the results were achieved and the level of confidence that the model has.

The goal of enabling explain ability in ML, as stated by “is to ensure that algorithmic decisions as well as any data driving those decisions can be explained to end-users and other stakeholders in non-technical terms”. S. Barocas, S. Friedler, M. Hardt, J. Kroll, S. Venka-Tasubramanian, and H. Wallach. The FAT-ML Workshop Series on Fairness, Accountability, and Transparency in Machine Learning. Accessed: July. 20, 2020. [Online]. Available: <http://www.fatml.org>

Explain ability sits at the intersection of transparency (consumers have the right to have decisions affecting them explained in understandable terms), causality (it is expected of the algorithms to provide not only inferences but also explanations), bias (the absence of bias should be guaranteed), fairness (it should be verified that decisions made by AI are fair) and safety (reliability of AI systems) (Hagras, 2018).

According to ((BSI), 2016), (Gasser & Almeida, 2017) we know that many machine learning algorithms have been labeled “black box” models because of their inscrutable inner-workings. What makes these models accurate is what makes their results difficult to interpret and understand they are very complex. So, even when some abstraction or transformation the models can be explainable, not always they are auditable. The discussion about audit AI is still open (Forum, 2019).

## 6. Augmented Intelligence at Telecom Argentina

More than a decade ago, in the former Cablevisión Argentina, we started working with Analytics to generate outcomes about our networks and services. Then, we began to innovate with other technologies related to Machine Learning and Artificial Intelligence and we continue after the 2018 merger, in Telecom Argentina.

Nowadays, as a STEM management our mission at Telecom Argentina is:

- Find new AI and ML based technologies to add value to the business.
- Define the technological strategy and emerging technologies.
- Develop innovative scientific tools to improve our infrastructure based on the demand for our services.
- Generate dimensioning models and performance parameters from statistical and mathematical analysis.

Our industry is undergoing a process of digital transformation and we, Telecom Argentina, are part of this process. Being part of this process also implies a transformation in our ways of working and in how we focus on them, but always oriented towards networks and services. So, we understood that we should jump to the next level.

We are not alone doing that; we belong to different Working Groups with the objective to define standards and best practices of these technologies. We also share our experiences and use cases.

They are:

- Artificial Intelligence and Machine Learning - SCTE
- Cross Industry AI/ML/Data Analytics Collaboration - CableLabs
- DOCSIS Data Analytics Work Group - CableLabs

Below, we detail some of the challenges we face from the emergence of this new AI that is human-machine collaboration.

### **6.1. VMAF: a tool for measuring video based on human perception**

VMAF (Video Multimethod Assessment Fusion) is a model proposed by Netflix in 2016 and is a video quality metric that combines human vision modeling with ML in order to provide a great viewing experience to their members.

VMAF is a fusion of elementary metrics into a final metric using a Support Vector Machines regressor which assigns weights to each elementary metric (Visual Information Fidelity, Detail Loss Metric, Temporal Information). In this way, the final metric is able to preserve all the strengths of the individual metrics and deliver a more accurate final score.

The model is trained and tested using a dataset of several 10 seconds long clips of different video genres and content characteristics. Those videos are distorted using different resolutions and bitrates. Then a subjective experiment is performed where a focus group of non-expert viewers score a source clip and the same clip distorted. From this experiment, the Mean Opinion Score (MOS) is calculated and it is also used to train the model.

Last year we have been working in an adaptation of this algorithm to our own OTT video platform: FLOW. We designed our datasets based on FLOW catalog and produced the distortions. Then, we performed some focus groups in order to obtain the corresponding MOS [ref SCTE 2019]. After doing that, we train and test the model with the elementary metrics mentioned before but applied to our content, the obtained MOS from the focus group and the SVM parameters that fitted the best.

With this tool we can qualify videos in a scale from “bad” to “excellent”, not only from an objective perspective but considering human perception too, in an automatic way with the help of machines. Previously, this task was performed by video experts, consuming a lot of time and resources. One of the advantages of VMAF is that video experts can dedicate time to analyze other video aspects and making decisions that machines can’t do and letting them do the repetitive tasks. This is a clear example of human-machine collaboration.

### **6.2. STEM-ML: a tool for capacity planning**

We developed a neural network-based machine learning tool for planning the capacity of the DOCSIS network [SCTE 2017], which was clearly superior to the tools used at the time.

However, it was very costly for us to adopt the tool by the teams responsible for capacity planning.

The business teams won't use the models because they do not know how they made their decisions and cannot be sure they work as advertised. It was difficult for us to explain to them that the decision to divide or not a node results from patterns in the data that the ML algorithm found and fixed into the model.

We finally understood that it was very risky for them to rely on models that cannot be explained.

From that moment on, we involve the teams from the moment they raise the need for an augmented intelligence tool. We focused on humans who for example carried out mobile network planning.

Although we have not yet used the formal tools that are being developed to explain the models, we have involved the team from the beginning of the development of STEM-ML for the planning of our 4G network. Applying agile methodologies for development. Explaining each model using and providing the planning team with the design criteria. Once an algorithmic model is up and running, the team must test the model carefully to see that it is operating in a reasonable way.

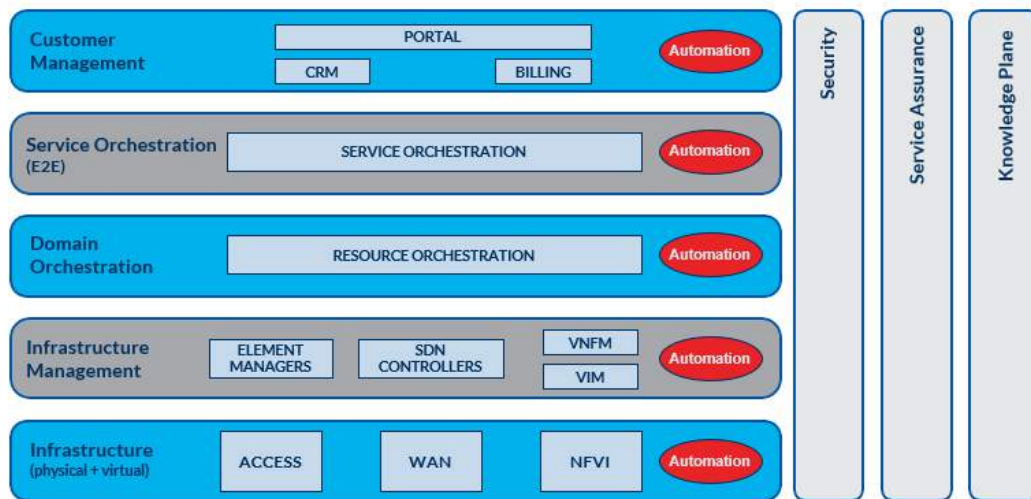
We were also developing with the team how to make the comparison between the results of the model and the methodology they had been using before.

We are defining with the team different levels of testing before putting the tool into production.

### **6.3. Telecom Argentina Knowledge Plane**

We at Telecom Argentina adopt the knowledge plane paradigm for our Telco Cloud project [paper 2019]. During the proofs of concept that we will begin to deploy with next year, we will not only advance on AI / ML technologies, but it is essential to understand this abstraction as an AgI tool that will allow us to operate our future networks and services.

At the beginning, and in order to have a common language, a framework was defined in which the preponderance of the knowledge plane in the different layers of the network, both virtual and physical, is highlighted.



**Figure 8 – Telecom Argentina Knowledge Plane**

There is no doubt that in the course of time this general vision will not only be modified but also specified. A result that only experience can provide.

## 7. Conclusion

Humans and machines working in collaboration can have a powerful impact on the effectiveness of business processes. Augmented Intelligence overcomes the limitations of isolating human understanding from the massive amounts of available data to analyze complexity in record time.

Telecom Argentina has begun the journey to AgI. In a jungle of recommendations from the different SDOs for each technology, the proof of concept with vendors and developers will introduce us to the best practices of human-machine collaboration, human beings contributing with interpretation, critical thinking and multiple perspectives or looks to the solutions that AI provides us.

## Abbreviations

3GPP	third generation partnership project
5G	fifth generation of mobile networks
5GCN	5G core network
AgI	augmented intelligence
AI	artificial intelligence
AMC	autonomic management and control
ARIMA	autoregressive integrated moving average
BBU	band base unit
BNG	broadband network gateways
CAPEX	capital expenditure
CDN	content delivery network
CLI	command line interface

CM	cable modem
CMTS	cable modem termination system
DCAE	data collection, analytics and events project
DE	decision element
DNS	domain name servers
DOCSIS	data over cable service interface specification
DSLAM	digital subscriber line access multiplexer
ENI	experiential networked intelligence
EPC	evolved packet core
ETSI	European Telecommunications Standards Institute
FCAPS	fault configuration accounting performance security
FM	failure management
FTTH	fiber to the home
GAN	generic autonomic networking architecture
GPON	gigabit-capable passive optical network
GSMA	global system for mobile communications association
HFC	hybrid fiber-coaxial
IBN	intent-based networking
IoT	internet of things
IPBB	ip backbone
IT	information technology
KDN	knowledge-defined network
KP	knowledge plane
MANO	management and orchestration
MDAF	management data analytics function
MIMO	multiple input multiple output
mIoT	massive internet of things
ML	machine learning
MNO	mobile network operator
MOS	mean opinion score
MSO	multi system operator
NaaS	network as a service
NE	network element
NFV	network functions virtualization
NOC	network operation center
NSSF	network slice selection function
NSSI-MDAF	network slice subnet instance management data analytics function
NWDAF	network data analytics function
NW-MADF	network management data analytics function
OLT	optical line termination
ONAP	open networking automation platform
ONT	optical network terminal
OPEX	operating expenditure
O-RAN	open radio access network
PAC	probably approximately correct
PCF	policy control function
PLMN	public land mobile network
PM	performance management

PNF	physical network functions
QoE	quality of experience
QoS	quality of service
RAN	radio access network
RAN-MDAF	radio access network management data analytics function
RAT	radio access technology
RIC	ran intelligence controller
SA	service and system aspects
SCTE	Society of Cable Telecommunications Engineers
SND	software defined network
SDO	standards developing organizations
SLA	service level agreement
SOC	service operation center
SON	self-organizing network
STEM	science, technology, engineering and mathematics
SVM	support vector machines
VIM	virtualized infrastructure manager
VMAF	video multimethod assessment fusion
VNF	virtual network functions
WG	working group
XAI	explained artificial intelligence
ZSM	zero-touch service and network management

# **Managing the Coronavirus Bandwidth Surge**

## **How to Cope with the Spikes and Long-term Growth**

A Technical Paper prepared for SCTE•ISBE by

**John Ulm**

Engineering Fellow, Broadband Systems  
CommScope – CTO Network Solutions team  
[john.ulm@commscope.com](mailto:john.ulm@commscope.com)

**Dr. Thomas Cloonan**

CTO – Network Solutions  
CommScope  
[tom.cloonan@commscope.com](mailto:tom.cloonan@commscope.com)



# Table of Contents

Title	Page Number
1. Introduction.....	4
2. Pre-Pandemic – the Calm before the Storm .....	4
3. COVID-19 Pandemic hits – the Storm Surge .....	6
3.1. NCTA COVID-19 Dashboard website .....	6
3.2. Comcast COVID-19 Dashboard website .....	9
3.3. Other data gathered from various MSOs around the Global .....	10
4. The Cause of the Surge – the usual Suspects .....	13
4.1. Work @ Home.....	13
4.2. Remote Learning.....	13
4.3. Video Streaming.....	13
4.4. Social Networking.....	13
4.5. Gaming .....	13
5. QoE Impacts of the Surge – the Good, the Bad & the Ugly .....	14
5.1. Upstream Capacity – Achille’s Heel of Cable .....	14
5.2. BW Impacts as measured by Speed Tests.....	14
5.3. Handling Congestion – DOCSIS to the Rescue.....	15
6. Repairing Leaks during the Storm – Short-term Fixes.....	17
6.1. Simple I-CCAP related changes.....	17
6.2. CM related changes .....	17
6.3. In-Home changes .....	18
6.4. Plant Optimizations.....	18
7. Cleaning up after the Storm – Near-term 6-12 month Strategies .....	18
7.1. CMTS related changes .....	18
7.2. Legacy Video.....	19
7.3. CM & In-Home changes .....	19
7.4. RF Plant Upgrades.....	20
8. Preparing for the next Storm – Network Capacity Planning.....	20
8.1. The “Basic” Traffic Engineering Formula.....	20
8.2. Providing Sufficient Headroom for a BW Surge.....	21
9. Building the Storm-proof Network – Mid- to Long-term Strategies .....	23
9.1. Mid-term Solutions – for 1-3+ Years of Capacity Needs .....	23
9.1.1. CCAP related changes.....	23
9.1.2. IP Video Migration.....	24
9.1.3. CM & In-Home changes .....	24
9.1.4. RF Plant Upgrades.....	24
9.2. Long-term Solutions – for 5-10 Year Horizon .....	25
10. Conclusion .....	25
Bibliography & References .....	27
Abbreviations.....	28

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Downstream Average Bandwidth per Subscriber through Jan ‘20 .....	5
Figure 2 – Upstream Average Bandwidth per Subscriber through Jan ‘20 .....	5
Figure 3 -Coronavirus BW Surge impacts across Time of Day .....	6
Figure 4 – NCTA COVID-19 Dashboard: National Peak Internet Growth .....	7
Figure 5 – NCTA COVID-19 Dashboard: Peak Usage status as of mid-May .....	8
Figure 6 – Comcast COVID-19 Bandwidth Chart.....	9
Figure 7 – Comcast COVID-19 Quotes .....	10
Figure 8 – Downstream Service Group Utilizations from No. American Metro-area.....	11
Figure 9 – Upstream Service Group Utilizations from No. American Metro-area .....	11
Figure 10 – Downstream Service Group Utilizations per Hour from N.A. Metro-area.....	12
Figure 11 – Upstream Service Group Utilizations per Hour from N.A. Metro-area .....	12
Figure 12 – Ookla Speedtest: U.S. Internet Performance during Pandemic .....	15
Figure 13 – Ookla Speedtest: U.S. Internet Latency during Pandemic.....	15
Figure 14 – Upstream Bandwidth depiction before the Coronavirus BW Surge.....	16
Figure 15 – Upstream Bandwidth depiction after the Coronavirus BW Surge.....	16
Figure 16 – QoE-based Traffic Eng Formula at Work during DS BW Surge, small SG .....	22
Figure 17 – QoE-based Traffic Eng Formula at Work during DS BW Surge, large SG .....	22
Figure 18 – QoE-based Traffic Eng Formula at Work for Marginal SG .....	23

# 1. Introduction

“With all due respect, sir, I believe this will be our finest hour.” That quote from the movie *Apollo 13* is strangely appropriate for what happened with the broadband industry after a global lockdown at the start of the COVID-19 pandemic. The Coronavirus pandemic turned the whole world upside down, with entire countries forcing their population to live, work, learn and play from home. There was a year’s worth of bandwidth growth (or more) in a 2-week interval. The broadband industry held its breath to see how our broadband infrastructure would cope.

Broadband kept society and the economy running (as best we can) in these hard times. This is arguably its most significant contribution to society in its short life-span. Hats off to everyone who has created and helped make broadband and the Internet available to society.

In this paper, the bandwidth (BW) impact to our broadband networks from around the globe is reviewed and showed how it was handled. The upstream got crunched much more than the downstream. Network capacity planning was key to having sufficient headroom to withstand an unexpected jolt. Some network capacity guidelines are reviewed and show how the Quality of Experience (QoE) level varies with different margins.

For those severely congested networks, some helpful tips when in crisis are provided. These overnight quick fixes can give a little breathing room until more permanent capacity can be added. The relative merits of some near-term solutions to deploy over coming months are discussed. These include: more DOCSIS channels, especially OFDM/OFDMA, more CCAP ports, segment congested nodes, rapidly increase DOCSIS 3.1 (D3.1) modem penetration, and deploying Wi-Fi 6 services.


Some bandwidth projections and mid- to long-term migration strategies are reviewed that operators should consider moving forward before the next bandwidth surge hits. These include: Migrate to 1218/85 or 1218/204 MHz plant today, with a transition to DOCSIS 4.0 over time; reduce legacy Video QAMs using IPTV/SDV/compression; Fiber Deeper; DAA; and Wi-Fi 6E. It is more important than ever to plan and start to implement our long-term strategies like 10G, especially for the upstream – cable’s Achilles heel.

Finally, the longer-term impact of the pandemic is considered. What is the new normal? We don’t expect bandwidth levels to ever go back to pre-COVID days. We’ll touch on how user’s behavior patterns changed and what new applications had the door opened.


## 2. Pre-Pandemic – the Calm before the Storm

Before we dive into the bandwidth impacts of the pandemic, let’s take a look at the world pre-COVID. CommScope/ARRIS has been monitoring subscriber usage for over a decade now from the same group of MSOs. The data from this set has been compared and maps closely to many other MSOs globally.

The chart below, Figure 1, shows the average subscriber downstream consumption, Downstream (DS) Average BW per subscriber (Tavg), during peak busy hours for a number of Multiple System Operators (MSOs) over a ten year period. At the start of 2020, DS Tavg had surpassed the 2 Mbps barrier. It turns out that the Tavg growth rate was higher at the start of this decade and has tailed off a bit in recent years. Over the last 3-4 years, this group of MSOs had an average downstream traffic growth that had been around 30%.



**Figure 1 – Downstream Average Bandwidth per Subscriber through Jan '20**



**Figure 2 – Upstream Average Bandwidth per Subscriber through Jan '20**


Interestingly, the upstream (US) traffic is growing at a significantly slower rate than the downstream as shown in Figure 2. During the same ten year period, the upstream Tavg generally grew at less than 20% compound annual growth rate (CAGR).

Over the years, traffic has also become more asymmetric with video applications driving downstream consumption [EMMEN\_2014]. The average DS:US ratio is about 14:1, the MSO with the largest DS:US ratio seems to have stabilized around an 18:1 ratio. The MSO with the lowest DS:US ratio was about 11:1 ratio.

### 3. COVID-19 Pandemic hits – the Storm Surge

During mid-March, our lives got turned upside down and the entire world suddenly had to work, learn and play from home. To monitor the impact of the Coronavirus BW surge, we had access to live data through the CommScope ServAssure program and were able to monitor 1000's of nodes to study the effects.

Below are some very representative data samples taken during late March to demonstrate some of the changes seen. Broadband usage started picking up earlier in the morning and slowing down later at night such that the lull in the wee hours of the morning shrunk. Daytime usage was up significantly. Some networks saw daytime downstream usage increase as much as 86% while daytime upstream usage increased as much as 150%. The shape of the upstream was also dramatically altered. The upstream would reach its peak around 10am and then stay plateaued for the rest of the day until about midnight.



**Figure 3 -Coronavirus BW Surge impacts across Time of Day**



While the staggering daytime usage increases made for good press, it did not impact the overall network capacity requirements. Networks must be designed for peak utilization. The daytime usage was starting from a much lower point, so its percentage increase is exaggerated. Our focus going forward will be on the increase to the peak busy period for the day. For the above nodes that we monitored, we saw peak downstream increases in the 12-25% range while peak upstream increases were even higher in the 20-50% range.

#### 3.1. NCTA COVID-19 Dashboard website

During the pandemic, the NCTA (The Internet & Television Association) was particularly proactive in publishing its members' network status. It collected nationwide information from the following MSOs:

- AlticeUSA, CableOne, Charter, Comcast, Cox, GCI, Mediacom, Midco, Sjoborg's

Altogether, this group represents 94% of the entire U.S.A. cable subscriber base. The NCTA published its data on the website: <https://www.ncta.com/COVIDdashboard>. The National Upstream Peak Growth and National Downstream Peak Growth charts are shown in figure 4 below.




**Figure 4 – NCTA COVID-19 Dashboard: National Peak Internet Growth**

These charts show the change in peak consumer usage relative to March 1<sup>st</sup>, before the pandemic lockdowns began. By early April, the upstream peak was up 35% while the downstream peak was up 20%. After flattening out for a bit, these peak increases started to gradually drop. As of mid-July, the upstream peak was still 24% higher than March 1<sup>st</sup>, while the downstream peak was down to just a 9% increase from the start. There has been a slight uptick with the end of summer.

At this point, it is not clear if the dip in the peak increases is a result of the gradual re-openings happening around the country, or whether this is just a typical summer slow-down as people start to do more outside

activities instead of being in front of their computers. It will also be interesting to observe the behavior of the above curves in the Autumn of 2020 with many schools planning to re-open with both face-to-face and on-line classes; this may result in an interesting bandwidth surge with more ubiquitous push towards on-line learning than was carried out in the Spring of 2020. Another interesting observation in the above curves is that the downstream peak increase of 9% after four months is actually in line with a 30% annual growth rate. So the downstream may be close to being back to normal levels. The upstream peak is definitely still elevated considering its historical 20% annual increases.

The NCTA website provides a wealth of additional information. It shows a dashboard with peak usage network status for the nation or for a given state. A snapshot of the National dashboard for mid-May is shown in figure 5. As can be seen, 97% of networks nationally were within their normal peak usage range with 2.8% at an elevated peak usage (i.e. yellow zone). Only 0.1% of networks were in the orange zone of substantially elevated peak usage and hardly any were in the red zone with significant elevated peak usage. This shows how well our broadband networks fared.




**Figure 5 – NCTA COVID-19 Dashboard: Peak Usage status as of mid-May**

The NCTA website also provided some additional key takeaways. Here were some given in mid-May:

- National US peak **growth remains mostly flat** with slight dip from 35% Peak
- National DS peak **growth receding** over last 2 months from 20% Peak
- Provider backbone networks **have significant capacity**
  - Show no signs of congestion
- US peak hours in many regions have shifted from late evening to afternoon
- Wi-Fi data traffic & Wi-Fi calling are increasing as compared to mobile
- Networks are supporting more Wi-Fi-connected devices

### 3.2. Comcast COVID-19 Dashboard website

During the first few months of the pandemic, Comcast also published information on their own website, <https://corporate.comcast.com/covid-19/network>. Figure 6 below shows the Comcast Bandwidth Chart from mid-May. The upstream peak growth was holding fairly steady around 30-32%, while the downstream peak growth shot up to ~20% in late March and then slid back down to ~10% in mid-May.



**Figure 6 – Comcast COVID-19 Bandwidth Chart**


Some observed effects on bandwidth as seen by Comcast include:

- Weekday usage is up:
  - VoIP & Video Conferencing is up 210-285%
  - VPN traffic is holding steady, up 30-40%
- Evening & weekend usage is up:
  - Gaming downloads are up 20-35% generally, up to 80% during new releases
  - 20-40% increase in streaming and web video consumption
  - Linear video consumption increased +2 hours per day per household
  - Video OnDemand (VoD) hitting record highs, up 50% YOY
- Xfinity Mobile sees a 36% increase in mobile data usage over Wi-Fi
  - But a 17% decline in LTE Data usage

Comcast also provided a number of interesting quotes, some of which are shown in figure 7. Of particular note is the quote from Comcast President Tony Werner: “We engineer the network to handle spikes and shifts in usage, and what we have seen so far with COVID-19 is within our capacity”. The authors have seen this numerous times –

*Networks that have been designed with sufficient capacity have been resilient and has handled the Coronavirus BW surge.*





**Figure 7 – Comcast COVID-19 Quotes**

### 3.3. Other data gathered from various MSOs around the Global


The CommScope ServAssure program allowed us to collect data from multiple MSOs from around the globe. We were able to look at data from before and after the start of the pandemic.

One set of data came from a North American metro-area and covered 1000’s of nodes. From that we were able to get the network utilization levels for every service group during the four weeks in March. The relative downstream (DS) service group (SG) utilization levels are shown in figure 8. The left side of the chart shows the entire node population while the right side zooms in on the “yellow” to “red” regions.


During the first week before the lockdown, almost 90% of the SGs were in the “green”, i.e. <65% DS utilization. These networks were in great shape and in no immediate need for additional capacity. Around 10% of SG were in the “yellow” or “orange” zones (i.e. between 65% to 80%) indicating that they might need extra DS capacity in the next 3-9 months. Only ~1% of SG were in the “red” that needed immediate attention (i.e. “red” means >80% utilization).

As utilization increased dramatically with the pandemic lockdown, there was a noticeable change in the DS SG utilization numbers by the last week in March. The number of DS SG in the “green” dropped to

about 80%, while DS SG in the “yellow” almost doubled to 18%. The DS SG in the “red” tripled to around 3%.



**Figure 8 – Downstream Service Group Utilizations from No. American Metro-area**



**Figure 9 – Upstream Service Group Utilizations from No. American Metro-area**


Figure 9 provides the upstream (US) service group (SG) utilization for the same North American metro-area. The entire node population is shown on the left while the right chart zooms in on the SG in the “yellow” to “red” zones.

This particular MSO has been meticulous with plant maintenance and reducing cascade lengths. Most of their upstream networks have at least four DOCSIS ATDMA channels. So in general, they have made sure there is plenty of upstream capacity. The early March data bares this out as approximately 98% of the US SG were in the “green”; ~1.5% were in the “yellow” and less than 0.25% of US SG were in the “red”.


As the pandemic hit, it took its toll. The “green” US SG dropped to almost 90%. The number of US SG in “yellow” quadrupled to ~6% while “red” US SG also quadrupled to 1%. While these percentages look low, this is across 1000’s of nodes that were impacted. The number of nodes that needed immediate attention increased dramatically, as well as the number of nodes needing extra capacity in the next 3-9 months.

Figure 10 takes a closer look at the downstream impacts using a typical DS SG. The chart on the left shows the utilization levels for different hours during the day and how it changed from week to week during March. The chart on the right shows similar utilization levels for the evening hours.

Notice how the DS utilization builds during the course of the day, from 10am to 2pm to 4pm to 6pm. By the 6-7pm hour, the DS SG is approaching its peak busy period. In general, 9-10pm was the busiest hour and used for our Tavg calculations. Although, 8-9pm and 10-11pm were very close. It shows that the DS peak busy period is close to a 3 hour window. There was a noticeable drop-off in utilization after 11pm.



**Figure 10 – Downstream Service Group Utilizations per Hour from N.A. Metro-area**



**Figure 11 – Upstream Service Group Utilizations per Hour from N.A. Metro-area**

Figure 11 takes a closer look at the upstream impacts using another typical US SG. The chart on the left shows the US utilization levels for different hours during the day and how it changed from week to week during March. The chart on the right shows similar US utilization levels for the evening hours.

Right away, the differences from the DS utilizations in figure 10 are obvious. By 10am hour, the US utilization is up to 65% and then the rest of the day it stays near peak levels. A closer look at the evening hours on the right shows how steady the US utilization remains for the entire evening. The peak busy period is now 11am to 11pm for this SG.

## **4. The Cause of the Surge – the usual Suspects**

Taking a closer look at the various applications that caused the Coronavirus BW surge, it turns out it is basically just a lot more of things that are already been seen.

### **4.1. Work @ Home**

The first very obvious impact of the lockdown was the significant increase in work@home employees. Some estimates showed that pre-COVID only 5%-10% of the workforce would work@home on a full or part-time basis. Overnight with the lockdown, this number skyrocketed to the 50%-90% range depending on locale. Working from home increases the use of Virtual Private Networks (VPN) and file transfers. Video conferencing would also jump up using programs such as WebEx, Microsoft Teams, Skype, Zoom, etc.

In general, many of these work applications are symmetric in nature, so they would tend to have a much bigger impact on the upstream capacity. Many current work environments are leveraging a cloud infrastructure and common applications such as Microsoft Office will routinely auto-save files to the cloud. This has big ramifications as employees work from home. If a worker is updating a 20MB PowerPoint or Word document, the cloud-sync function may cause it to get uploaded a dozen times over an hour of working. Again, adding additional burden to our upstream networks.

### **4.2. Remote Learning**

Schools were one of the first institutions to close down when the pandemic hit, and school districts scrambled to set up their remote learning programs. Zoom and Skype became two of the favorite video conferencing methods to reach students during the day. Again, a symmetric application putting additional burdens on the upstream.

YouTube usage also saw a big uptick as a way for teachers to distribute their lesson plans. And with everyone at home, YouTube became a big source of educational material as people researched the pandemic and looked for ways to pass the time.

### **4.3. Video Streaming**

As was noted by Comcast above, viewers had more time on their hands and hence the viewing hours per week went up. Many video content providers offered free trial subscriptions to entice folks to sample their content. It remains to be seen how many of these consumers will stick with the service after the trial is over.

### **4.4. Social Networking**

With chatting around the water cooler removed, there was a definite increase in social networking apps as folks looked at ways to stay in touch. Facetime, Skype and Zoom sessions became very popular, not only to keep in touch with current contacts, but to re-establish with some long lost folks too.

### **4.5. Gaming**

With everyone at home and more spare time available, gaming applications such as Twitch also a sparked increase. For many of today's on-line games, a download might be on the order of 100GB. There were several instances of a new release of a popular game that directly impacted network performance.

While, in general, gaming applications tend to impact the downstream more, they may often be accompanied by some social networking interaction and possibly some video conferencing.

The gamers are the ones that may have been impacted the most as these networks became congested and latencies increased.

## **5. QoE Impacts of the Surge – the Good, the Bad & the Ugly**

With network congestion comes impacts on subscribers Quality of Experience (QoE). The Coronavirus BW surge introduced 12-18 months' worth of upstream BW growth in a 1 week period. The downstream was only slightly more forgiving with 8-12 months of BW growth when the lockdowns started.

Some of the negative QoE effects on subscribers include:

- Packet delays for gamers from higher latencies and buffer-bloat
- Video tiling in streaming & video conferencing from packet delays and drops
- File Download & Upload times increased (a little bit)

In addition to the above impacts, some subscribers who are working/learning/playing from home have insufficient service tiers to support their traffic load and/or are running up against any data caps that their service provider may have.

### **5.1. Upstream Capacity – Achilles' Heel of Cable**

And this heavy usage is lasting longer throughout the day. The upstream peak period now lasts 12-14 hours from mid-morning to midnight. Because the upstream pipe is much smaller while having the more significant BW surge, it has come under the most scrutiny. At this point, the upstream is considered the Achilles' Heel of the cable industry. Current North American 42 MHz systems might support 60-100 Mbps from 3-4 DOCSIS 3.0 ATDMA channels. European 65 MHz HFC systems have a little more breathing room with up to 175 Mbps from up to 7 DOCSIS 3.0 ATDMA channels. But this capacity pales compared to the 5+ Gbps of downstream capacity found on an 870 MHz plant. Because of this, a lot of the focus on following sections will be on enhancing the upstream capacity.


### **5.2. BW Impacts as measured by Speed Tests**

For many consumers, when network performance starts to lag, they run a speed test. One of the most common available is Ookla's [www.speedtest.net.Ookla](http://www.speedtest.net.Ookla). Ookla tracks and publishes its speed test results from around the globe. They did a very good job of showing how every country fared during the pandemic:


<https://www.speedtest.net/insights/blog/tracking-covid-19-impact-global-internet-performance/>

Figure 12 shows a snapshot of the USA Internet performance based on the Ookla speed test. It covers a timeframe from December 2019 to May 2020. When the lockdowns first occurred in mid-March, the Fixed broadband networks recorded a 9% drop. Note that this was actually less than the drop seen during the holiday week of December 23-29.

The downstream internet performance then gradually recovered over the following weeks until it was only down about 1% from March 2<sup>nd</sup> until mid-May.



**Figure 12 – Ookla Speedtest: U.S. Internet Performance during Pandemic**



**Figure 13 – Ookla Speedtest: U.S. Internet Latency during Pandemic**

Figure 13 shows the latency measurements from Ookla speed test during this same period. Notice that there was no perceived impact on latencies due to the Coronavirus BW surge. More evidence that broadband networks are holding their own!

### 5.3. Handling Congestion – DOCSIS to the Rescue

Overall, most of these bad QoE effect have been relatively minimal to date for most service groups. Only a handful of SGs that were pushed to the limit before the BW surge had detrimental QoE. In general, the DOCSIS networks are holding up to this sudden, stressful packet load. Why, despite heavily congested upstreams causing delays and drops? Several reasons:

- Good network capacity planning added plenty of headroom
  - e.g. enough headroom to absorb the largest Service Level Agreement (SLA) bursts
- Excellent CMTS Scheduling algorithms
  - These are sophisticated AI Engines that have evolved over the last two decades
  - They are excellent at adapting to congestion, with fair BW distribution

- Most internet applications are elastic and forgiving; TCP and ABR recover from throughput reductions & packet loss


And it also doesn't hurt that subscribers are more tolerant about small lapses when they are more worried about the virus et. al.

To bring home this point of having good bandwidth capacity planning with an excellent CMTS scheduler, figures 14 and 15 show a before and after picture of a hypothetical upstream channel. These upstreams will have traffic from many different service flows (SF) from many different modems. These will present a fairly static load on the system. But at any instant in time, any one of these SF might burst up to its maximum SLA capacity. This is shown in the yellow spikes in the figures.

Before the BW surge with good network planning, then the available US capacity is sufficient to handle the largest SLA burst. Everyone is happy!



**Figure 14 – Upstream Bandwidth depiction before the Coronavirus BW Surge**



**Figure 15 – Upstream Bandwidth depiction after the Coronavirus BW Surge**

After the Coronavirus BW surge, many of the service flows increase their traffic load and the tide rises. There is no longer sufficient room to handle the largest BW bursts. This is shown in the red circles in the figure. This is also when the CMTS schedulers kick into high gear.

Great CMTS schedulers will throttle down these high bursts, delaying packets and perhaps dropping packets when needed. They are also looking across all SFs and prioritizing as needed. The bottom line is that these bursts get spread out slightly in time with barely a notice from the subscriber. Many subscribers with smaller SLA's may not see any impacts at all. This is the beauty of statistical multiplexing traffic and intelligent US scheduling during the Coronavirus BW surge.

## 6. Repairing Leaks during the Storm – Short-term Fixes

For those handful of SGs that became super-congested during the BW surge, operators were looking for some quick fixes that might gain an additional 5% or 10% or 15% of capacity gain. They also wanted something that they could implement in the matter of hours or days to relieve the immediate heartburn.

Below is a smattering of small things that might help. By themselves they may not add up too much, but collectively a handful of these might be enough to get an operator through the rough patch until more permanent capacity can be added to the system.

### 6.1. Simple I-CCAP related changes

Today's Integrated CCAP (I-CCAP) are very powerful and flexible networking devices with a lot of capabilities. The most straightforward solution to the BW crunch is to add more DOCSIS channels which can be done under software (SW) control using licensing. This would be the first choice solution provided there is spectrum available.

If the operator has a substantial mix of DOCSIS 3.0 and 3.1 modems, then they will need to consider whether to add 3.0 SC-QAM channels or 3.1 OFDM/A channels. In general, the 3.1 OFDMA US channels provide 65% to 100% better spectral capacity (Mbps/MHz) than 3.0 ATDMA channels. The 3.1 OFDMA channels are also more robust and can operate over a much wider frequency range, even well below 20 MHz in the “muck”. In general, the 3.1 OFDMA channels have been under-utilized by many operators to date.

But even some 3.0 ATDMA channels could be operated below 20 MHz as well. The operator might need to use a lower modulation (e.g. 16-QAM or QPSK) and/or lower channel widths with increased FEC and interleaving. Some I-CCAP might have an integrated US agility capability to auto-adjust QAM levels.

If there is not sufficient available spectrum in the downstream, the 3.1 OFDM channel is robust enough to run in the roll-off region. E.g. put an OFDM channel from 750-942 MHz in a 750 MHz system. There might be 1 Gbps or more of capacity available in the roll-off.

In a very congested network, the CMTS scheduler is using the SF Tmax value (i.e. its SLA burst value) to fairly assign bandwidth. So, for example, maybe all SF will receive 60% of Tmax during congestion. But there might be a big discrepancy between the highest and lowest Tmax SLA. For example, the top tier in the upstream might be 50 Mbps while the lowest tier is 5 Mbps. The lowest tier will only get 10% of the capacity of the high tier. One strategy might be to temporarily boost the lowest tier (e.g. up to 20 Mbps) to ensure that these customers get adequate US performance to work successfully from home.

The DOCSIS CMTS has dozens of other configuration parameters that could be tweaked to help things. One could enable “Power-Boost” modes; configure Tmin parameters to provide minimum BW guarantees; adjust SF priorities; or tweak any scheduler parameters that might be available.

### 6.2. CM related changes

There are a handful of cable modem (CM) based configuration items that can be done to help. First, the operator should turn on the TCP Ack suppression on all modems if they haven't already. This will help reduce the number of TCP Acks in the upstream and every little bit helps when severely congested.

Buffer-bloat is another problem that has been studied for many years. This is the large increase in queue depths during congested intervals. This then results in very high latencies seen by the applications. The



operator can enable some Buffer Control TLVs and turn on DOCSIS 3.1 Active Queue Management (AQM) to help suppress buffer-bloat and reduce latencies when congested.

### **6.3. In-Home changes**

From the subscriber's home perspective, there are a number of items that fall under the quick fix umbrella. First, an operator might encourage subscriber to upgrade to a higher service tier. A new promotion that is geared towards families that are working at home and remote learning might gain a lot of traction while bringing in extra revenue.

With so many people in the home trying to actively use the internet simultaneously, many Wi-Fi routers may be caving under the load. An operator might provide some guidelines to their customers on improving the Wi-Fi router location within the house.

Finally, some I-CCAP might have capabilities like Integrated Service Class Agility (ISCA) that can dynamically reduce the Tmax value temporarily for any "bad actors".

### **6.4. Plant Optimizations**

Sometimes operators will use a common configuration across all of their downstream and upstream CMTS ports. This may result in using the lowest common denominator and many SG may be leaving unused capacity on the table. For heavily congested SG, the operator may need to optimize the configuration parameters for those channels.

Depending on the Signal-to-Noise Ratio (SNR), an operator may be able to increase the QAM modulation and/or optimize the Forward Error Correction (FEC) for a given channel. For example on a cleaner plant, the US ATDMA channel might get bumped from 32- to 64-QAM modulation. An US OFDMA channel might go from 256- to 512- or even 1024-QAM modulation. In the DS OFDM channel, an operator might jump from 1024- to 4096-QAM.

The I-CCAP may ship with a default FEC setting that is very robust, but not efficient. Optimizing the US FEC could increase US capacity by as much as 13%. That may be enough to get that SG out of the dog house.

**Warning** – only change the QAM-modulation and FEC to values that can be supported by the SNR. Being too aggressive may introduced undetected packet errors causing dropped packets and making performance worse than before. An operator should use tools like the Intelligent Channel Optimizer (ICO) to optimize settings. For OFDMA channels, optimize bit-loading settings in Modulation Profiles using PMA tools (while still maintaining adequate Packet Error Rate, PER).

## **7. Cleaning up after the Storm – Near-term 6-12 month Strategies**

The previous section talked about some quick fixes that could be done to relieve pressure on SGs in the "red" zone. This section takes a look at what can be implemented over the next 6-12 months to address SGs in the "yellow" zone. These can prepare an operator in case another BW surge arises.

### **7.1. CMTS related changes**

Rinse and repeat – add more DOCSIS channels is the recurring theme here. Any additional capacity means adding more DOCSIS channels. The questions become: where does the spectrum come from and what is the mix of 3.0 and 3.1 channels.

Once the available spectrum is used including the roll-off region, then the operator will need to look at reclaiming spectrum from Legacy video QAMs; upgrading plants to 1218/85 or 1218/204 MHz; and/or segmenting service groups in either hubs or in nodes. Segmenting SG implies that the operator will have to add more CCAP ports to their system.

These options will be detailed below.

## **7.2. Legacy Video**

For many operators, Legacy video still consumes a significant portion of their downstream spectrum. This is predominantly broadcast video QAM channels but also may include a handful of Video-on-Demand (VOD) QAMs.

Operators should consider converting any remaining older MPEG-2 encoded broadcast programs over to MPEG-4 encoding. This cuts their bandwidth requirements in half. This also means they will need to retire any MPEG-2 only settop boxes (STBs) that are still out there.

With all the improvements in encoding and stat-muxing technology, the operator may be able to increase the number of programs per QAM channel. For example, they may be able to jump from 2-3 HD programs per QAM up to 4-5 HD programs per QAM.

Perhaps the most powerful tool available to the operator for reducing Legacy video spectrum is Switched Digital Video (SDV). This technology has been around for 15 years but has undergone a bit of a renaissance by migrating to cloud-based infrastructures. This mature technology is now more cost effective than ever. SDV plays well with I-CCAP that integrate the SDV video QAMs directly into it. You can turn on an SDV system in a matter of weeks. And by aligning the SDV group with the shrinking DOCSIS SG, the operator could free up enough spectrum for a pair of 192 MHz OFDM channels.

## **7.3. CM & In-Home changes**

The operator's biggest weapon in its arsenal is DOCSIS 3.1. But it doesn't help if there are not enough 3.1 modems in the field. The first priority for the operator should be to upgrade ALL modems for high tier customers (e.g. 200Mbps+ DS). By getting all the high tier customers onto 3.1 OFDM/OFDMA channels, they get better service. And this off-loads the potentially congested 3.0 SC-QAM channels for the lower tiers.

As mentioned before, the upstream is the cable operator's Achilles heel. 3.1 OFDMA channels can more than double the upstream capacity in a 42MHz system. Operators should be rolling out more and more OFDMA channels. As the 3.1 modem penetration starts to overtake 3.0 penetration, then 3.0 SC-QAM channels can be reclaimed and converted to OFDMA spectrum.

Another big in-home improvement that an operator can make in the near-term is to start upgrading their customers to the new Wi-Fi 6 routers (i.e. 802.11ax). These can significantly increase the in-home Wi-Fi capacities and reach, which are both critical with so many people working from home.

The Wi-Fi 6 routers could be bundled with 3.1 modems for new service tiers that target work and learn at home families.

## 7.4. RF Plant Upgrades

There are a number of things that an operator can do with respect to the outside plant to enhance capacity in the near term. Perhaps the lowest hanging fruit with minimal investment is to segment existing nodes. Typically nodes are deployed with a single downstream module and a single upstream module. This configuration is referred to as 1x1 and maps to a single CMTS US port and a single DS port. The node may have two or four RF ports, and this DS and US are shared across all RF ports.

Some nodes with a digital return can easily be re-configured to support a second US. This effectively splits the node into 1 DS and 2 US or a 1x2 node. This will help relieve US congestion on that SG.

Later, another DS module can be added to the node to make it a 2x2 configuration. Eventually more modules can be added so the node evolves into a 4x4 configuration. All of this can be done at the node location without touching the rest of the outside plant. Additional CCAP ports will be needed in the hub site as well.

For severely congested SG, an operator might consider upgrading those first to 1218/85 or 1218/204 MHz plants. This is part of the longer term strategy and provides enough capacity for the rest of the decade. The 85 MHz return almost triples the amount of usable US spectrum. Combined with 3.1 OFDMA, it can support 400-500 Mbps US tiers. The 204 MHz return enables 1 Gbps US services. Upgrading all of your plants to 1218 MHz will take a while, so start today with those SG that are most congested.

DOCSIS 3.1 is very powerful and optimizes capacity according to the quality of the plant. It turns out that long fiber runs (e.g. >40km) with AM optics can have a detrimental effect on 3.1 OFDM capacity. This can be remedied by replacing those AM fiber runs with digital optics when Remote PHY devices (RPD) or Remote MAC-PHY devices (RMD) are placed in the node. And again, this might be part of a longer term MSO strategy to migrate to a complete Distributed Access Architecture (DAA).

## 8. Preparing for the next Storm – Network Capacity Planning

The CommScope (formerly ARRIS) team has been providing industry leading research in traffic engineering for many years which was most recently highlighted in [ULM\_2019]. Some additional references of note include [CLO\_2014], [EMM\_2014], [ULM\_2014], [CLO\_2016], [ULM\_2016], [ULM\_2017] and [CLO\_2017].

### 8.1. The “Basic” Traffic Engineering Formula

Previously, [CLO\_2014] provided an introduction to traffic engineering and quality of experience (QoE) for broadband networks. From there, the paper went on to develop a relatively simple traffic engineering formula for cable service groups that is easy to understand and useful for demonstrating basic network capacity components.

The “Basic” formula shown below is a simple two-term equation. The first term ( $N_{sub} \cdot T_{avg}$ ) allocates bandwidth capacity to ensure that the aggregate average bandwidth generated by the  $N_{sub}$  subscribers can be adequately carried by the service group’s bandwidth capacity. The first term is viewed as the “DC component” of traffic that tends to exist as a continuous flow of traffic during the peak busy period.

#### The “2014” Traffic Engineering Formula (Based on $T_{max\_max}$ ):

$$C \geq (N_{sub} \cdot T_{avg}) + (K \cdot T_{max\_max}) \quad (1)$$

where:

C is the required bandwidth capacity for the service group

Nsub is the total number of subscribers within the service group

Tavg is the average bandwidth consumed by a subscriber during the busy-hour

K is the QoE constant (larger values of K yield higher QoE levels)...

where  $0 \leq K \leq \text{infinity}$ , but typically  $1.0 \leq K \leq 1.2$

Tmax\_max is the highest Service Tier (i.e. Tmax) offered by the MSO

The MSO data collected and shown in section 2 provides a good indication of Tavg, at least before the Coronavirus BW surge hit. Note that figures 1 and 2 are very generalized results that are averaged across millions of subscribers. A given operator will need to take a look at their own networks to ascertain what is an appropriate value for them. Regarding SG congestion, they need only look at the total SG consumption (i.e. Nsub\*Tavg product term) and there is no need to break it down into the individual Nsub and Tavg components.

There are obviously fluctuations that will occur (i.e. the “AC component” of traffic) which can force the instantaneous traffic levels to both fall below and rise above the DC traffic level. The second term ( $K \cdot \text{Tmax\_max}$ ) is added to increase the probability that all subscribers, including those with the highest Service tiers (i.e. Tmax values), will experience good QoE levels for most of the fluctuations that go above the DC traffic level.

The second term in the formula ( $K \cdot \text{Tmax\_max}$ ) has an adjustable parameter defined by the K value. This parameter allows the MSO to increase the K value and add bandwidth capacity headroom that helps provide better QoE to their subscribers within a service group. In addition, the entire second term is scaled to be proportional to the Tmax\_max value, which is the maximum Tmax value that is being offered to subscribers. A change in the K value results in a corresponding change within the QoE levels experienced by the subscribers who are sharing the service group bandwidth capacity (C). Lower K values yield lower QoE levels, and higher K values yield higher QoE levels).

In previous papers [CLOONAN\_2013, EMM\_2014], found that a K value of ~1.0 would yield acceptable and adequate QoE results. [CLOONAN\_2014] goes on to provide simulation results that showed a value between  $K=1.0$  and  $1.2$  would provide good QoE results for a service group of 250 subscribers. Larger SGs would need larger values of K while very small SGs might use a K value near or less than 1.0.

## 8.2. Providing Sufficient Headroom for a BW Surge

It will be useful to provide some examples on how the CommScope Basic Traffic Engineering formula can be applied to provide adequate headroom to sustain a BW surge. The first DS example is shown in figure 16. It assumes that the SG size is 100 subs. Before the BW surge (on the left in the figure), it uses a Tavg value = 2.36 Mbps taken from figure 1. The highest tier is 1G (i.e. Tmax\_max = 1 Gbps) with a QoE constant  $K = 1.2$ . The basic formula says that the operator needs at least 1436Mbps to maintain that Kvalue. This might be a combination of 3.0 SC-QAM and 3.1 OFDM channels bonded together. The Nsub\*Tavg component equals 236 Mbps (shown in blue) with the  $K \cdot \text{Tmax\_max}$  component = 1200 Mbps (shown in yellow).

After the surge is shown on the right hand side of figure 16. It assumes DS usage grows by 25% to Tavg = 2.95 Mbps. The Nsub\*Tavg has now increased to 295 Mbps while overall capacity has remained fixed at 1436 Mbps. This reduces the  $K \cdot \text{Tmax\_max}$  component to 1141 Mbps which is an effective K value

equal to 1.14. This is still in the “good” QoE range, despite seeing an instantaneous 25% increase in usage.

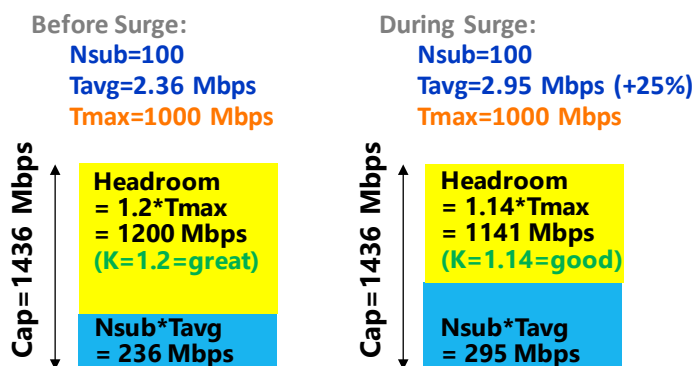


Figure 16 – QoE-based Traffic Eng Formula at Work during DS BW Surge, small SG

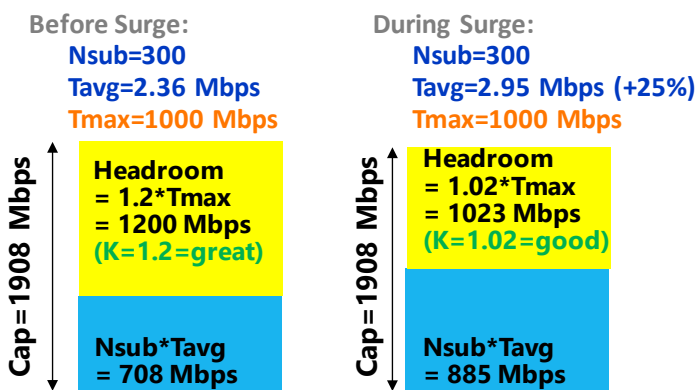


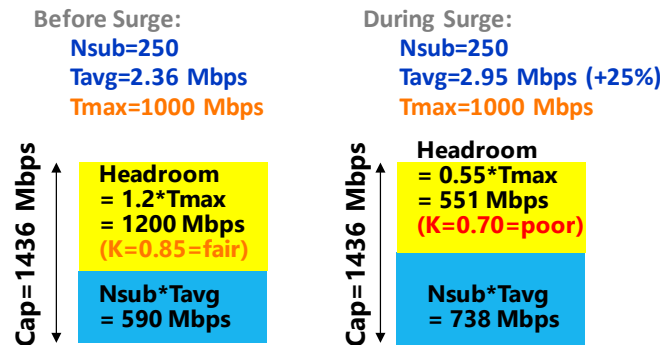
Figure 17 – QoE-based Traffic Eng Formula at Work during DS BW Surge, large SG

The second example in figure 17 increases the number of subscribers in the SG to 300 subs. Before the surge, this increases the  $N_{sub} * T_{avg}$  component to 708 Mbps. The overall required capacity is now increased up to 1908 Mbps to still maintain a QoE constant  $K = 1.2$ .

During the surge, the increased  $T_{avg}$  now has a bigger impact due to the increased SG size. The  $N_{sub} * T_{avg}$  component has now swelled to 885 Mbps. This now leaves 1023 Mbps for the  $K * T_{max\_max}$  component which provides an effective QoE constant  $K = 1.02$ . This is still in the “good” region, although near the bottom. So for larger SG, it is more important to have the cushion associated with  $K=1.2$ .

The third example in figure 18 takes a look at what happens if sufficient headroom is not applied. This example assumes the operator has rolled out the same 1436 Mbps overall capacity across their entire footprint, regardless of the SG size. This example looks at a SG with 250 subs. The  $N_{sub} * T_{avg}$  component before the surge is 590 Mbps. This makes for an effective  $K = 0.85$  which is in the “fair” region. This means that the 1G customers will have good QoE much of the time, but may see some congestion periodically during peak busy periods.

Once the Coronavirus BW surge hits this SG, its  $N_{sub} \cdot T_{avg}$  jumps to 738 Mbps. This reduces the  $K \cdot T_{max\_max}$  component down to 551 Mbps which results in  $K=0.7$ . This is in the “poor” region. The 1G customers QoE impact becomes more noticeable.



**Figure 18 – QoE-based Traffic Eng Formula at Work for Marginal SG**

It is important that operators not only monitor their network utilizations, but they must also understand the highest tier being offered and adjust their thresholds accordingly. For example, just having a utilization threshold of 80% to take action is not good enough by itself. That might have been adequate when the top tier was 100 Mbps, but then starts to fail when the top tier is 1 Gbps.

*As the top tier is increased, the utilization threshold to take action must be correspondingly decreased to maintain consistent QoE.*

## 9. Building the Storm-proof Network – Mid- to Long-term Strategies

The Coronavirus BW surge is a glimpse at our future BW trends. We will not return to our pre-COVID days and will need to deal with a “new normal”. This surge should serve as a wake-up call to make sure our industry continues to push its networks forward. We’ll take a look at some mid-term solutions that can be implemented over the next couple years and then look at what will be coming later this decade.

### 9.1. Mid-term Solutions – for 1-3+ Years of Capacity Needs

Over the next couple years, operators need to focus on getting switched over to DOCSIS 3.1. This is the industry’s workhorse right now. Operators should have plans in place to try and become 100% 3.1, phasing out their 2.0 and 3.0 modems. This can be done over a 5-8 year window, but needs to start in earnest now. Remember, DOCSIS 3.1 is your biggest weapon with your biggest bang for the buck.

#### 9.1.1. CCAP related changes

As time goes on, operators should continue to maximize their 3.1 OFDM/OFDMA capacity. Any older CMTS should be upgraded to support the full 3.1 capabilities, including OFDMA in the upstream with both 85 and 204 MHz splits; and multiple 192 MHz OFDM channels in the downstream.

Over time, SG sizes will continue to shrink with node segmentations, so new CCAP ports will continually be added. Segmenting SG size is especially helpful in containing  $T_{avg}$  growth as was shown in the examples of the previous section.

During this timeframe as 3.1 penetration increases and high tier customers are moved to 3.1 modems, the operator should start considering switching some 3.0 SC-QAM channels to OFDM/OFDMA spectrum.

### **9.1.2. IP Video Migration**

As part of the strategy to completely remove the Legacy Broadcast video spectrum, operators should have their IP video migration plans in full swing. It may take 5 to 8 years or more to get all of the Legacy STB swapped out for IP STB or other IPTV devices.

Ramping up the IPTV consumption while it coexists with Legacy Broadcast QAMs may create a sizable bandwidth bubble. The operator may also need to consider SDV or encoding enhancements as described in section 7.2.

Some other technologies to consider when deploying Adaptive Bit-Rate (ABR) IP video is multicast-assisted ABR (M-ABR) and Smart ABR technologies. Using multicast in an IP video world may still be very important when considering special events with large viewer populations tuned in. M-ABR also significantly reduces the BW requirements between the CDN and the CCAP. For a CCAP with 25,000 subscribers, it could potentially reduce IP video capacity originating at the CDN from several hundred Gbps down to 10's of Gbps.

### **9.1.3. CM & In-Home changes**

During the mid-term phase, operators should get the remaining high tier subs completely on 3.1 modems. Not only does this make it easier to provide them with good QoE during times of congestion, it will free up precious 3.0 bandwidth that will be needed by the lower service tiers.

And with higher penetrations of 3.1 modems, the operator can make better use of OFDMA spectrum. OFDMA has the potential of increasing the 42 MHz upstream capacity from 100 Mbps with 3.0 to the 200-250 Mbps range depending on plant quality.

In addition to the migration to 3.1 modems, operators should also be deploying Wi-Fi 6 throughout the rest of its customer base.

### **9.1.4. RF Plant Upgrades**

During the mid-term years, operators should be upgrading their older 550/750/870 MHz plants to 1218/85 or 1218/204 MHz plants. These technologies are available today and will carry the operator through the rest of this decade. [ULM\_2019] takes a look at how 1218 MHz networks can achieve our Cable 10G DS goals. DOCSIS 3.1 specification was written in 2014 and is just now becoming widespread. DOCSIS 4.0 specification just became available and requires an even bigger plant investment with new taps. So don't expect 4.0 to become mainstream until the end of this decade. And if growth rates continue to slow, maybe the need for 4.0 will be pushed out further in time.

The long term strategy for every cable operator should be to push fiber deeper. Eventually we may see fiber passing every home, but that level of investment takes multiple decades. However, MSOs should continue the fiber deep march in the mid-term and continue to reduce serving area size and cascade lengths.

There are some significant benefits with a fiber deep approach (e.g. N+0, N+1). In addition to the reduced SG size and increased 1.2GHz spectrum, the total number of active components can be reduced

increasing reliability AND cutting OPEX costs in half. Fiber deep networks are also much closer to homes and become a jumping off point to selectively offer FTTH services.

Finally, some operators may be on the road to DAA. The mid-term period will be a good time to start that transition and start ramping up their deployments of RPD or RMD in the field.

## **9.2. Long-term Solutions – for 5-10 Year Horizon**

What does the distant future hold? Nobody knows for sure and the crystal ball is cloudier than ever. Bandwidth usage may not snap back to its original pre-Coronavirus levels when Coronavirus ends. Why? Because this novel social experiment we are all involved in may foster a new social paradigm. Workers and companies may decide to explore more work-at-home activities and students and schools may decide to explore more on-line education activities.

Bandwidth will obviously continue to grow into the 2020's, although the rate of growth has a lot of debate. We are starting to see new applications like eSports getting a boost right now. The problem experienced during Coronavirus is only a sampling of what will happen in the future when that bandwidth growth crosses certain thresholds

***MSOs & Vendors need to begin working now to upgrade the HFC Plant to support the future Bandwidth Growth.***

While DOCSIS 3.1 can carry us a very long way, the industry needs to work on new technologies today to be ready for the future. This is what the Cable 10G initiative and DOCSIS 4.0 specification are all about. DOCSIS 4.0 will give us a choice of two options. Full-Duplex (FDX) technologies for N+0 plants that support dynamically shared spectrum from 108-684 MHz. This will enable 4-5 Gbps US tiers.

Alternatively, 4.0 also specifies some ultra-high split options paired with extended spectrum (ESD). These can offer service tiers similar to FDX on a 1794/492 MHz plant using simple FDD techniques. Both of these technologies are enabling multi-gigabit upstream tiers to be competitive with 10G PON symmetric offerings.

Inside the home, there will also be a migration from today's Wi-Fi 6 routers to the next generation of Wi-Fi 6E. The new generation will have a new spectral band from 5.9-7.1 GHz that enables new use cases and opens consideration of 10 Gbps around the home wirelessly that pairs nicely with 4.0 modems.

## **10. Conclusion**

The Coronavirus BW surge gave us 12-18 months of upstream usage growth overnight, and about half that in the downstream. And despite the surge, we have seen that the DOCSIS network continues to work very well. But this event has shown weaknesses in the system that need filling to make it through the near-, mid- and long-terms.

Some operators needed some immediate quick fixes for severely congested service groups to relieve the pressure. These needed to be rolled out in a matter of hours or days. Adding DOCSIS channels with the I-CCAP was first choice if spectrum was available, with the roll-off being a good option for the 3.1 OFDM channels. Turning on OFDMA in the upstream should be a priority as well. There are a number of various configuration changes that could be done to improve operation including:

- Enabling TCP Ack suppression
- Reduce buffer-bloat latencies with Buffer-control TLVs, 3.1 AQM



- Optimizing QAM-modulation or FEC
- Re-locating in-home Wi-Fi

These might only provide 5% or 10% or 15% improvements, but everything helps in a congested network.

In the near-term over the coming months, operators can add more substantive improvements to their network capacity. To find spectrum for additional DOCSIS channels, operators can use SDV or improved encoding to reduce Legacy Video spectrum. Node segmentation might split SG to help manage BW. Nodes might go from 1x1 to 1x2 to 2x2/2x4 to 4x4 segmentation. This can be done simply by upgrading the node and not touching anything else in the outside plant. Meanwhile, more DOCSIS 3.1 modems and Wi-Fi 6 routers should be deployed to customers' homes.

Network capacity planning is key to being able to withstand another bandwidth surge. The CommScope/ARRIS Basic Traffic Engineering formula provides guidelines on how much capacity is required to provide good QoE. We saw some examples on what happens before and after the BW surge and why the QoE headroom is important.

*Networks that have been designed with sufficient capacity have been resilient and has handled the Coronavirus BW surge.*

In the mid-term over the next several years, many plants will be upgraded to 1218/85 or 1218/204 MHz. Fiber Deep networks will continue to be our long term goal as we look to push fiber closer to the home while increasing plant reliability and reducing OPEX. Distributed Access Architectures including RPD and RMD will take hold, especially on longer fiber runs where they provide the most performance gains. And the IP video migration should be in full swing with technologies like M-ABR supplementing it.

Looking into the distant future, BW will not snap back after the pandemic. It will continue to grow through the decade. MSOs need to begin upgrading their networks for the future. While DOCSIS 3.1 continues to be the workhorse today, development of 4.0 products will help to enable symmetric multi-gbps services later in this decade.

Broadband kept society and the economy running, as best we can, in these hard times. This is arguably its most significant contribution to society in its short life-span. Hats off to everyone who has created and helped make the Internet and Broadband available to society.

While collectively we should all be proud of the infrastructure we have created, we still have much work to do. The Coronavirus BW surge has exposed cracks in the system. This should serve as a wake-up call. We still have improvements to make to the infrastructure. It's time to start upgrading the network for the demands of the 2020s.

# Bibliography & References

[CLO\_2019] T. J. Cloonan et. al., “Capacity Planning, Traffic Engineering, And HFC Plant Evolution For The Next 25 Years,” SCTE Cable-Tec 2019, SCTE

[CLO\_2017] T. J. Cloonan et. al., “The Big Network Changes Coming with 1+ Gbps Service Environments of the Future,” SCTE Cable-Tec 2017, SCTE

[CLO\_2016] T. J. Cloonan et. al., “Using DOCSIS to Meet the Larger BW Demand of the 2020 Decade and Beyond,” NCTA Spring Technical Forum 2016, NCTA

[EMM\_2014] “Nielson’s Law vs. Nielson TV Viewership for Network Capacity Planning,” Mike Emmendorfer, Tom Cloonan; The NCTA Cable Show Spring Technical Forum, April, 2014

[FDX\_PHY] “DOCSIS 4.0 Physical Layer Specification”, CM-SP-PHYv4.0-D01-190628, Cablelabs 2019

[FDX\_XSD\_IBC] “Full duplex DOCSIS & Extended Spectrum DOCSIS Hold Hands to Form the 10G Cable Network of the Future”, by F. O’Keeffe et. al., IBC 2019

[RAM\_2020] Ram Ranganathan et. Al., “Decoding the Bandwidth Surge during COVID-19 Pandemic – an Indepth Study on DOCSIS Upstream Bandwidth Surge and their Impact on Video Conferencing”, SCTE Cable-Tec 2020, SCTE

[ULM\_2019] J. Ulm, T. J. Cloonan, “The Broadband Network Evolution continues – How do we get to Cable 10G?”, SCTE Cable-Tec 2019, SCTE

[ULM\_2018] J. Ulm, “Making room for D3.1 & FDX – Leveraging Something Old that is New Again!”, SCTE Journal of Network Operations : Find Fresh Approaches to Plant-Related Topics, Vol 4. No. 1. Dec 2018, SCTE

[ULM\_2017] J. Ulm, T. J. Cloonan, “Traffic Engineering in a Fiber Deep Gigabit World”, SCTE Cable-Tec 2019, SCTE

[ULM\_2016] “Adding the Right Amount of Fiber to Your HFC Diet: A Case Study on HFC to FTTx Migration Strategies”, John Ulm, Zoran Maricevic; 2016 SCTE Cable-Tec Expo

[ULM\_2014] “Is Nielsen Ready to Retire? Latest Developments in Bandwidth Capacity Planning”, John Ulm, T. Cloonan, M. Emmendorfer, J. Finkelstein, JP Fioroni; 2014 SCTE Cable-Tec Expo

# Abbreviations

ABR	Adaptive Bit Rate
AI	Artificial Intelligence
AQM	Active Queue Management
ATDMA	Advanced Time Division Multiplex Access
BAU	Business as Usual
Bps	Bits Per Second
BW	Bandwidth
CAA	Centralized Access Architecture
CAGR	Compounded Annual Growth Rate
CAPEX	Capital Expense
CCAP	Converged Cable Access Platform
CDN	Content Distribution Network
CM	Cable Modem
CMTS	Cable Modem Termination System
CPE	Consumer Premise Equipment
D3.1	Data Over Cable Service Interface Specification 3.1
DAA	Distributed Access Architecture
DEPI	Downstream External PHY Interface
DOCSIS	Data Over Cable Service Interface Specification
D3.1	DOCSIS revision 3.1
DS	Downstream
EPON	Ethernet Passive Optical Network (aka GE-PON)
EQAM	Edge Quadrature Amplitude Modulator
ESD	Extended Spectrum DOCSIS
FD	Fiber Deep
FDD	Frequency Division Duplex
FDX	Full Duplex (i.e. DOCSIS)
FEC	Forward Error Correction
FTTH	Fiber to the Home
FTTLA	Fiber to the Last Active
FTTP	Fiber to the Premise
FTTT	Fiber to the Tap
FTTx	Fiber to the 'x' where 'x' can be any of the above
Gbps	Gigabits Per Second
GHz	Gigahertz
HFC	Hybrid Fiber-Coax
HD	High Definition
HP	Homes Passed
HSD	High Speed Data
I-CCAP	Integrated Converged Cable Access Platform
ICO	Intelligent Channel Optimizer
IEEE	Institute of Electrical and Electronics Engineers
IPTV	Internet Protocol Television
ISCA	Integrated Service Class Agility

M-ABR	Multicast-assisted Adaptive Bit Rate
MAC	Media Access Control interface
MACPHY	DCA instantiation that places both MAC & PHY in the Node
Mbps	Mega Bits Per Second
MDU	Multiple Dwelling Unit
MHz	Megahertz
MSO	Multiple System Operator
N+0	Node+0 actives
NCTA	The Internet & Television Association
NFV	Network Function Virtualization
NSI	Network Side Interface
Nsub	Number of subscribers
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiplexing Access (Upstream)
OPEX	Operating Expense
OTT	Over the Top
PER	Packet Error Rate
PHY	Physical interface
PMA	Profile Management Application
PNM	Proactive Network Maintenance
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
QoS	Quality of Service
RF	Radio frequency
R-OLT	Remote OLT
RPD	Remote PHY Device
R-PHY	Remote PHY
R-MACPHY	Remote MAC-PHY
RMD	Remote MAC-PHY Device
RX	Receive
SCTE	Society of Cable Telecommunications Engineers
SDV	Switched Digital Video
SF	Service Flow
SG	Service Group
SLA	Service Level Agreement
SNR	Signal to Noise Ratio
STB	Set-top box
SW	Software
Tavg	Average bandwidth per subscriber
TCP	Transmission Control Protocol
Tmax	Maximum Sustained Traffic Rate – DOCSIS Service Flow parameter
TX	Transmit
US	Upstream
VOD	Video on demand
VoIP	Voice over Internet Protocol
VPN	Virtual Private Networks

# Low Latency Docsis: Concepts And Experiments

A Technical Paper prepared for SCTE•ISBE by

**Tushar Mathur**

Staff Systems Engineer, CTO Office  
CommScope Inc.  
90 Matheson Blvd. W., Mississauga, ON, Canada  
tushar.mathur@commscope.com

**Ram Ranganathan**

Director of Systems Engineering, CTO Office  
CommScope Inc.  
90 Matheson Blvd. W., Mississauga, ON, Canada  
ram.ranganathan@commscope.com

**Greg Gohman**, CommScope Inc.

**Bob Zhang**, University of Waterloo

# Table of Contents

Title	Page Number
1. Introduction.....	3
2. Low Latency DOCSIS Architecture And Goals .....	3
3. Experimental Setup .....	7
4. Experimental Analysis Of Latency In DOCSIS 3.1 System .....	8
4.1. Single Upstream Service Flow Experiments.....	9
4.1.1. Comparing Single Service Flow With LLD ASF Classic Service Flow .....	11
4.2. LLD ASF With Best Effort Scheduling Experiments.....	12
4.2.1. Scenario 1 .....	13
4.2.2. Scenario 2 .....	13
4.2.3. Scenario 3 .....	13
4.2.4. Scenario 4 .....	14
4.2.5. Summary Of LLD ASF Experiments .....	14
5. Conclusions And Future Work .....	16
Abbreviations .....	17
Bibliography & References.....	17

## List of Figures

Title	Page Number
Figure 1 LLD Architecture .....	5
Figure 2 Traditional DOCSIS REQ-GNT Cycle.....	6
Figure 3 PGS Granting Mechanism .....	6
Figure 4 Differentiated Services Byte in IP Header .....	7
Figure 5 Experimental Test Setup .....	8
Figure 6 Single SF BE Mean and 99 Percentile Of Gaming Traffic.....	11
Figure 7 Single SF BE 99 Percentile and Jitter Of Gaming Traffic.....	11
Figure 8 Single SF vs LLD ASF CL SF Mean And 99 Percentile Of Gaming Traffic.....	12
Figure 9 Single SF vs LLD ASF CL SF 99 Percentile and Jitter Of Gaming Traffic .....	12
Figure 10 LLD ASF Both SF BE Mean And 99 Percentile Of Gaming Traffic .....	15
Figure 11 LLD ASF Both SF BE 99 Percentile And Jitter Of Gaming Traffic .....	15
Figure 12 LLD ASF LL SF=PGS And CL SF=BE Mean And 99 Percentile Of Gaming Traffic .....	15
Figure 13 LLD ASF LL SF=PGS And CL SF=BE 99 Percentile And Jitter Of Gaming Traffic .....	16

## List of Tables

Title	Page Number
Table 1 Summary Of Number Of Traffic Pattern Streams Per Scenario .....	10
Table 2 Single Service Flow Gaming Traffic Latency With Best Effort and Proactive Grant Scheduling...	10
Table 3 LLD ASF Scenario 1 Results .....	13
Table 4 LLD ASF Scenario 2 Results .....	13
Table 5 LLD ASF Scenario 3 Results .....	14
Table 6 LLD ASF Scenario 4 Results .....	14

## 1. Introduction

Today's internet traffic typically comprises data, voice, or video traffic with no extraordinary means to logically segregate traffic based on its network latency sensitivity. Applications have varying requirements for bandwidth, latency or jitter. Some apps require high bandwidth, such as large file downloads or video traffic, and certain apps require low latency, such as online gaming traffic or high frequency trading. The online gaming industry is on a rapid growth path and has become an exciting mainstream revenue source. With an increasing demographic that streams gameplays, the cloud gaming services are bringing new online gaming experience closer to the consumers and it will require support from the 10G initiative driven MSOs to deliver the best quality of experience by ensuring *high* bandwidth and *low* latency or *low* lag support. By enhancing the user experience, the MSOs have an opportunity to generate a new revenue stream. Welcome to the world of Low Latency DOCSIS!

The LLD architecture as proposed by CableLabs enables a logical separation of the latency sensitive non-queue building traffic and regular queue-building internet traffic in to two separate queues. The two queues, Low Latency SF and Classic SF are encapsulated in an Aggregate Service Flow (ASF) to shape the traffic. A key innovation that is part of the LLD architecture is a new scheduling service known as Proactive Grant Scheduling (PGS) [1].

There are multiple sources of latency in DOCSIS networks, including protocol/application dependent queuing delays, propagation delay, Request-Grant delay, channel configuration (OFDM or SC-QAM interleavers, cyclic prefix, FEC, etc.), and switching/forwarding delays. The purpose of LLD is to reduce latency from two of these sources – protocol/application dependent queuing delays and Request-Grant delays.

This paper will focus on the LLD architecture basics and experimental results from the lab studies using the concept of an LLD ASF and PGS in the *DOCSIS Upstream*. The paper will also compare LLD capable system latency with classic latency.

## 2. Low Latency DOCSIS Architecture And Goals

Often times, the bandwidth or speed of a connection is confused with latency of an application. Bandwidth or speed of a connection means how much of the data can be downloaded or uploaded within a time interval. For example, to watch a 4K YouTube video or to download a Call Of Duty 25 GB game patch, a user would need a good bandwidth. There are times when the bandwidth of a connection isn't enough to deliver the best QoE. For example, a multiplayer game of Call Of Duty requires players to shoot at other players as well as download any real time rendering of dynamic game environment. This typically results in packets being transmitted at a bit rate of 100 kbps to 200 kbps in the Upstream and Downstream direction. It is important that the packets reach their destination as quickly as possible so that the player does not get shot themselves first and the game environment rendering is synced with a player's action. If the gaming environment actions are not synced with a players action then it is because of a "high lag" in the network. This time duration of the packets to reach the Call Of Duty gaming server and returning a response to the multiplayer gamer is called Latency. So, to deliver the best QoE it is important to maintain reasonable bandwidth and latency. Inside a home there are multiple users transmitting traffic in the US and DS. Some of the internet traffic may be file download, or a YouTube video, or a Netflix video and other may be gaming traffic or video conferencing.

Typically all the traffic will flow in to a single DOCSIS service flow, with a mix of traffic that builds queues like the video streaming apps and other traffic that doesn't build queues like a multiplayer gaming app. The problem with this architecture is that the gaming app gets treated similar to a video streaming app, appending non-queue building traffic in to queue building traffic. Hence, this adds latency & jitter to an already latency sensitive application.

The LLD architecture uses a logical construct called an Aggregate Service Flow (ASF) that encapsulates two underlying service flows – one for Non-Queue building traffic and the other one for Queue building traffic. The intention of separating application's traffic in to two logical queues is to make sure that the application data that builds queues in the DOCSIS access network don't cause delays for data that does not build queues.

For example, let us assume there are two applications that are transmitting traffic in the DOCSIS channel.

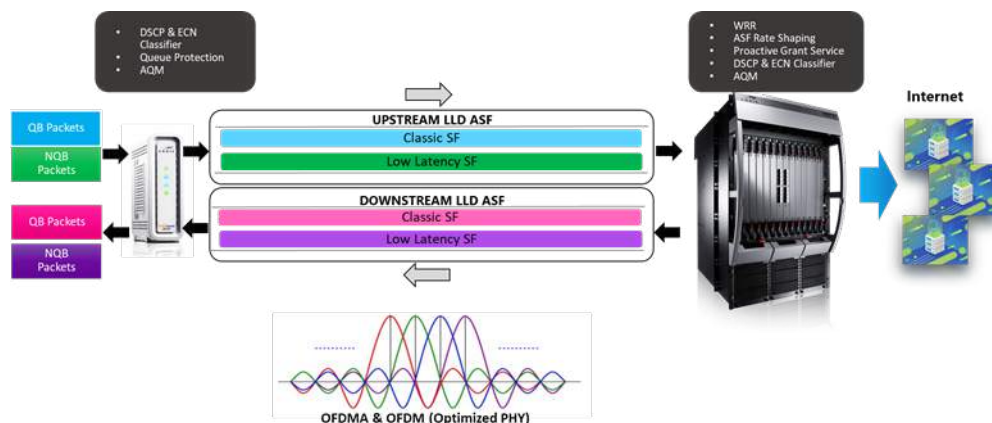
1. Online gaming application traffic is typically a few hundred kbps of UDP payload and will not cause queues to build in a SF. UDP does not have any congestion algorithms and will transmit at a set rate, without retransmissions. We call this traffic as the NQB or Low Latency traffic
2. Large file upload traffic in the order of a few MB of TCP payload and will cause formation of queues in a SF. The inherent nature of TCP congestion algorithms is to seek as much bandwidth as possible, cause retransmissions in case of packet loss which results in formation of queues. We call this as QB or Classic Traffic

As compared to a large file upload, the gaming traffic requires the best latency possible so that a gamer can have the best Quality of Experience. If both the traffic types i.e. gaming & file upload are contending in the same SF, although the UDP traffic is so low in bandwidth, it gets congested with the TCP traffic that causes queues to build-up. In other words, the a latency sensitive application gets delayed in the SF queue by a queue building application. So, by using LLD, the latency sensitive NQB application gets its own SF queue without impacting the latency of other QB applications.

Now, let's take a closer look at the features provided by the LLD architecture. Figure 1 shows the LLD architecture and its components. The CMTS has many important functions in defining the QoS in LLD:

1. A new ASF encapsulating LL and CL SF, known as the LLD ASF
2. A new Weighted Round Robin Scheduler for the two SFs
3. Rate shaping these two SFs at an aggregate level
4. A new scheduling type known as the Proactive Grant Service
5. Traffic Classification in to SFs using DSCP and ECN fields in the IP header
6. Active Queue Management – a new AQM algorithm called the Immediate AQM and Coupled AQM for the two constituent Service Flow
7. Queue Protection
8. Latency Histograms






**Figure 1 LLD Architecture**

The LLD ASF provides an encapsulation to the traffic shaping of the LL SF and CL SF by enforcing an Aggregate Maximum Sustained Rate (AMSR). LL SF and CL SF are not like traditional service flows that are shaped by an MSR value. There are always only two SFs in the LLD ASF. For example, if the ASF AMSR is set to 100 Mbps then the traffic flowing in to LL and CL SF can be 20 Mbps and 80 Mbps or it can be 30 Mbps and 70 Mbps respectively, or in any other proportion but bounded by the ASF AMSR. In the upstream LLD ASF, the traffic flowing in the individual SFs will depend on the number of grants each service flow receives.


The Granting mechanism in LLD is governed by a Weighted Round Robin Inter-SF Scheduler running on the CMTS. The scheduler is to behave with conditional priority of providing LL SF grant priority without starving the CL SF. The WRR weight is a configurable parameter with a maximum value of 255. A 230 value set in the LL SF would mean 90% weight (230/255) of grants is to be provided to the LL SF. The weights provided by scheduler does not result in an unfairness between the two service flows because of the Coupled AQM feature of the LLD architecture. More on that later in this section.

Traditionally, the DOCSIS upstream data transmission follows the mechanism described in Figure 2. As soon as a packet arrives, a Bandwidth (BW) Request (REQ) is transmitted to the CMTS to allocate bandwidth requested by the CM. CMTS responds to the CM by sending a Bandwidth Grant (GNT) to the modem based on QoS parameters governed by the CMTS. Once the modem receives the Bandwidth Grant in a MAP packet, it will process the MAP and transmit the data packet to the CMTS, from where it is routed towards a server in the internet. The time for the entire process is given by  $\Delta 1 + \Delta 2 + \Delta 3$ , where  $\Delta 1$  = Packet Processing Delay at the CM + Waiting for a REQ transmission slot + US propagation delay of BW REQ;  $\Delta 2$  = BW REQ processing delay at the CMTS + MAP generation delay + DS propagation delay + Wait Time for GNT;  $\Delta 3$  = US propagation delay of actual data to be transmitted




**Figure 2 Traditional DOCSIS REQ-GNT Cycle**

LLD introduces a new data scheduling type known as the Proactive Grant Service or PGS. PGS enables a faster request grant cycle by eliminating the need for a BW REQ even though it's not prohibited to send a BW REQ on a PGS SF. CM will still be able to send BW REQ if the SF bandwidth demands are not met by PGS alone. In PGS, BW GNTs are continuously sent to the modem at a Guaranteed Grant Rate (GGR in bps) and at a Guaranteed Grant Interval (GGI in microseconds) value as soon as some activity is detected by the CMTS's proprietary Activity Detection algorithm. GGI is the interval between successive data transmission opportunities. CMTS can track the bandwidth utilization and can adjust the GGR depending on any anticipated future demands, but at the time of writing this paper, the current CMTS implementation does not adjust GGR. If there is no activity detected on the SF then PGS will switch to sending unicast request opportunities at Guaranteed Request Interval (GRI in microseconds). The CMTS traffic shaper makes sure that the SF is not getting more grants than its maximum sustained rate by verifying the bounds of GGR and GGI. GGR, GGI, and GRI are configurable values in a PGS enabled flow. Figure 3 shows that with PGS, the upstream transmission time is shortened to  $\Delta 2' + \Delta 3$  where  $\Delta 2'$  includes DS propagation delay and a reduced processing time for MAP generation. It is typically less than or equal to the GGI. If GGI is set to a value less than  $\Delta 1 + \Delta 2$ , PGS will provide a reduction in the upstream transmission time.



**Figure 3 PGS Granting Mechanism**

In LLD, packet classification plays an important role in placing a packet into a particular SF – CL or LL. The classifiers in LLD examine the 8 bit Differentiated Services field in the IP header for DSCP (MSB 6 bits) and ECN value (LSB 2 bits). For example, the classification can be made based on a packet's DSCP field marked as EF or ECN field is set to ECT(1) or CE (see figure 4) then it will get mapped to the LL SF and any other traffic will be transmitted in the CL SF by default.



**Figure 4 Differentiated Services Byte in IP Header**

The Active Queue Management algorithms run on CL and LL SFs, both. The CL SF is to use the DOCSIS PIE algorithm, that drops packets as the queues build, to maintain a target latency defined by the configuration. The LL SF will implement a new AQM algorithm known as Immediate AQM, that does not drop packets but marks them with ECN bits. Default marking starts at 475  $\mu$ s and will always mark the packets beyond 1 ms of latency. As stated earlier, the AQMs act as coupled AQM on both of the constituent SFs i.e. the IAQM of LL SF is coupled to the DOCSIS PIE of CL SF. The coupling will act as a backstop on the LL SF if the CL SF is getting overwhelmed by traffic. When the CL SF throughput is overwhelming the system with requests for the grants to keep up, the ECN marking is induced in the LL flow. The induced ECN marking in the LL SF reduces the bandwidth demands in the LL flow and the remainder grants from the ASF token bucket will be available for the CL flow.

Queue Protection categorizes packets into the application data flows, termed Microflows. All packets of each Microflow are characterized by identical values in a set of header fields. QP algorithm must act on every Microflow that becomes a queuing source in the LL SF. If the LL SF buffers are getting filled at a critical threshold then that queuing source needs to be redirected in to the CL flow.

### 3. Experimental Setup

There are some delays associated with the Upstream and Downstream PHY layer of DOCSIS. These delays can be minimized depending on an MSOs network conditions because there are trade-offs between channel robustness vs latency. For the experiments in this paper, the following PHY and MAC layer settings were configured.

In the upstream, an OFDMA channel with symbols per frame (k) set to 16 with 50 kHz subcarriers. The Cyclic Prefix value = 0.9375  $\mu$ s and the Rolloff Period = 0.3125  $\mu$ s were set for the channel to reduce PHY latency. The channel width of the OFDMA channel was set to 42 MHz. With 1K QAM, this resulted in a channel capacity of 336.60 Mbps.

In the downstream, a 192 MHz wide OFDM channel was configured. Cyclic Prefix = 0.9375  $\mu$ s and Rolloff Period = 0.625  $\mu$ s. The time interleaver depth = 1 (Depth of the time interleaver in symbols of an OFDM channel). With 4K QAM, the resulting downstream channel capacity is approximately 2100 Mbps.


There are some delays associated with the MAC layer of DOCSIS. These delays can be minimized depending on an MSO's network conditions. Values of these parameters can be adjusted to further reduce DOCSIS MAC latency.

MAP-size of the channel can be adjusted between 1 to 13. MAP-size is configured as an average size in 800 microseconds ticks. A setting of map-size = 1 implies 800  $\mu$ s. Lowering the MAP-size results in faster acquisition of a bandwidth slot. For these experiments, the MAP-size was set to 2 i.e. 1600  $\mu$ s based on a MSO feedback. More MAP means messages mean more opportunities to transmit in the upstream. There is a tradeoff between choosing lower MAP-size with downstream bandwidth. Lower the MAP-size, more the downstream bandwidth consumed by Maps.

"max-round-trip-delay" is the RF RTT from a cable plant that can be configured in the CMTS. This delay should be adjusted based on expected distance between the Upstream burst receiver and the CM. In our experiments the value was set to 800 microseconds in propagation delay that equals to 100 miles in distance between the CMTS and CM.

Databackoff configuration can assist in spreading the effect of collisions in the broadcast request opportunities in a highly bursty and congested upstream. In this experimental setup we set it to a value in the range 5-8.

Figure 5 shows a high-level network diagram of a complex test rig to measure end-to-end DOCSIS Upstream latency. To emulate a typical MSO service tier, the CM's SLA was set to 100 Mbps in the DS and 20 Mbps in the US.



**Figure 5 Experimental Test Setup**

For all the experiments that were done using the PGS scheduling, the parameters were set as GGR = 2 Mbps, GGI = 1080  $\mu$ s, GRI = 540  $\mu$ s. The Weighted Round Robin scheduler is set to default weight of 9:1 ratio for the constituent service flows. The IAQM algorithm in LL SF is using default parameters of Ramp Function Exponent = 19 and the threshold = 1000  $\mu$ s. The DOCSIS PIE AQM algorithm in CL SF is using a target latency = 25 ms based on MULPI spec recommendation to set AQM latency target between the range of 10 ms to 100 ms. The service flow buffer size is set to 50 ms based on previous experimentation and recommendations.

## 4. Experimental Analysis Of Latency In DOCSIS 3.1 System

Since the LLD architecture is defined in the realm of the DOCSIS 3.1 standard, we will first dive into an experimental analysis of latency in DOCSIS 3.1 system. CableLabs has used simulation studies to analyze the dual-queue architecture. While those studies were useful, real-world empirical studies using

real-world non-deterministic data traffic model, real-world CCAP with real-world Schedulers and Mappers will provide even more valuable information to MSOs as they try to deploy LLD in the field.

Before we start our deep dive in to the experiments, it is important to note that all the information is experimental since a lot of MULPI spec specific development is still under development on the CMTS and CM software.

## **4.1. Single Upstream Service Flow Experiments**

The first set of experiments used a DOCSIS 3.1 modem with AQM enabled on a single Service Flow emulating a single home with multiple users. This a traditional Service Flow setup that ingests all types of traffic i.e. QB and NQB. The AQM in this scenario is set to a target latency of 25 ms and uses DOCSIS PIE algorithm. In this experiment, we monitored the gaming traffic latency. One gaming stream was always transmitted in the LL SF and another in CL SF, in the case of LLD ASF experiments in order to measure the impact of QB and NQB on the gaming traffic.

### **Scenario 1**

This scenario acts as a baseline experiment. The traffic mix included 2 UDP gaming streams in the US and DS directions, and a simple web browsing session.

### **Scenario 2**

In this scenario, the traffic mix included 2 UDP gaming streams in the US and DS directions, 2 web browsing sessions, 2 ABR video streams (DASH) sessions that emulate OTT content like Netflix or YouTube, and an Upload speed test.

### **Scenario 3**

In this scenario, the traffic mix included 2 UDP gaming streams in the US and DS directions, 2 web browsing sessions, 2 ABR video streams (DASH) sessions that emulate OTT content, and a file upload session emulating picture or short video upload.

### **Scenario 4**

In this scenario, the traffic mix included 2 UDP gaming streams in the US and DS directions, 2 web browsing sessions, 2 ABR video streams (DASH) sessions that emulate OTT content, and a file upload session emulating picture or short video upload on a social media platform, and two video conferencing sessions.

Table 1 provides a summary of number of streams of different traffic mixes for all of the four Scenarios.

**Table 1 Summary Of Number Of Traffic Pattern Streams Per Scenario**








							
Scenario 1	1x	1x	1x				
Scenario 2	1x	1x	2x	2x	1x		
Scenario 3	1x	1x	2x	2x		1x	
Scenario 4	1x	1x	2x	2x		1x	2x

Table 2 covers all 4 scenarios that are mentioned above for single service flow setup.

The first scenario helped baseline the behavior of BE vs PGS for low-bit rate gaming traffic. The PGS enabled service flow showed a mean latency of ~1.5 ms and jitter of ~0.3 ms for the gaming traffic stream. Meanwhile, a BE enabled service flow showed a mean latency of ~5.5 ms and jitter of ~0.8 ms for the gaming traffic. This baseline behavior shows that use of PGS reduced latency by ~72%!

We can see that a single US SF with BE scheduling type has high latency (mean, 95<sup>th</sup> percentile, and 99<sup>th</sup> percentile) as the traffic mixes are increased in different scenarios. It is important to note that jitter\* in all tests except Scenario 1 is pretty high. If there is anything that impacts a gamers QoE more than latency, it is jitter.

In all the experiments of this paper, Jitter is measured as the Mean Absolute Packet Delay Variation.

**Table 2 Single Service Flow Gaming Traffic Latency With Best Effort and Proactive Grant Scheduling**

Scenario	Scheduling Type	Mean Latency (ms)	95 Percentile Latency (ms)	99 Percentile Latency (ms)	Jitter* (ms)	US Throughput (Mbps)	DS Throughput (Mbps)	PGS Efficiency (%)
<b>1 (Baseline)</b>	BE	5.5	7.7	7.9	0.8	2	2	N/A
<b>1 (Baseline)</b>	PGS	1.5	1.7	2.7	0.3	2	2	44
<b>2</b>	BE	27	37	42	79	20	35	N/A
<b>3</b>	BE	27	43	50	142	20	35	N/A
<b>4</b>	BE	27	47	54	148	20	36	N/A

\*Mean Absolute Packet Delay Variation




Figure 6 Single SF BE Mean and 99 Percentile Of Gaming Traffic




Figure 7 Single SF BE 99 Percentile and Jitter Of Gaming Traffic

#### 4.1.1. Comparing Single Service Flow With LLD ASF Classic Service Flow


So how does the Single Service Flow latencies compare to the LLD ASF Classic SF latencies? In a Single Service Flow, QB and NQB traffic fill the SF buffers together, while in a LLD ASF there is a separation of QB and NQB traffic. We want to compare the latency metrics of gaming traffic in a Single Service Flow versus the latency metrics of gaming traffic within the Classic Service Flow that carries the QB traffic. The expectation in this experiment is that the latency metrics of gaming stream in a Single Service Flow will be approximately equal to gaming stream latency metrics within the Classic Service Flow.

Based on the data in Table 2 and Figure 8, it can be observed that the gaming traffic mean latency in CL SF is under the Single SF Mean up to Scenario 3, but for Scenario 4 the values are at a close approximation. These data points confirm that the theoretical expectation is correct.

The 99 percentile level for CL and single SF is approximately the same. This should be a good indicator that the NQB traffic within QB traffic in LLD architecture is fairly treated like the current circumstances of using a single service flow. In other words, latency performance isn't a zero-sum game. By separating




QB and NQB traffic, the NQB traffic can achieve much better latency performance, without degrading the performance of the QB traffic.



**Figure 8 Single SF vs LLD ASF CL SF Mean And 99 Percentile Of Gaming Traffic**

Figure 9 shows the 99 Percentile and Jitter data for the same comparison, and it can be observed that the CL Jitter is tracking to the single SF Jitter.



**Figure 9 Single SF vs LLD ASF CL SF 99 Percentile and Jitter Of Gaming Traffic**

## 4.2. LLD ASF With Best Effort Scheduling Experiments

This next set of experiments were focused on using LLD ASF, where a separation of Queue Building traffic and Non-Queue Building traffic was done by using relevant classifiers. By default, Queue Protection was enabled in these experiments and so was AQM as described in the experimental setup section.

The goal of this section is to highlight the importance of separating QB and NQB traffic. All the experiments were done with two configurations – 1. LL SF with BE scheduling and 2. LL SF with PGS scheduling. This is to contrast the behavior of two scheduling types.



### 4.2.1. Scenario 1

Table 3 shows the results of Scenario 1. The first scenario acts as our baseline for the separation of QB and NQB traffic since it includes only two gaming streams, one in each flow and the web session traffic in CL SF. It was observed that gaming latency is the same when the Best Effort scheduling is used in both the constituent SFs, CL and LL SF. Gaming traffic latency in LL SF is reduced by ~72% when PGS scheduling is used in the LL SF. The jitter also drops by ~70% when PGS is used in the LL SF.

**Table 3 LLD ASF Scenario 1 Results**

ASF With BE In LL SF and CL SF									
Scenario	Flow Type	Scheduling Type	Mean Latency (ms)	95 Percentile Latency (ms)	99 Percentile Latency (ms)	Jitter (ms)	US Throughput (Mbps)	DS Throughput (Mbps)	PGS Efficiency (%)
1.00	Low Latency	BE	5.33	6.81	6.97	0.93	1.00	3.00	N/A
1.00	Classic	BE	5.35	6.77	6.97	0.93	1.00		N/A
ASF With PGS in LL SF and BE in CL SF									
1.00	Low Latency	PGS	1.27	2.28	2.50	0.29	1.00	3.00	0.44
1.00	Classic	BE	5.24	6.68	6.89	0.86	1.00		N/A

### 4.2.2. Scenario 2

In Scenario 2, the traffic is increased significantly. The impact of QB and NQB is much clearer for either of the tests – with BE in both SF vs. with PGS in LL & BE in CL SF. The gaming traffic in LL SF continues to experience low latency and jitter compared to QB and NQB traffic mix in the CL SF.

**Table 4 LLD ASF Scenario 2 Results**

ASF With BE In LL SF and CL SF									
Scenario	Flow Type	Scheduling Type	Mean Latency (ms)	95 Percentile Latency (ms)	99 Percentile Latency (ms)	Jitter (ms)	US Throughput (Mbps)	DS Throughput (Mbps)	PGS Efficiency (%)
2.00	Low Latency	BE	5.54	7.13	7.37	0.88	1.00	35.00	N/A
2.00	Classic	BE	22.93	35.09	37.71	124.29	19.00		N/A
ASF With PGS in LL SF and BE in CL SF									
2.00	Low Latency	PGS	1.19	1.62	1.79	0.10	1.00	35.00	44.00
2.00	Classic	BE	23.41	33.51	37.29	104.34	19.00		N/A

### 4.2.3. Scenario 3

In Scenario 3, the traffic is more bursty because of frequent FTP file upload sessions. The Jitter experienced in CL SF is upwards of 150 ms which is detrimental to gamers QoE. The queues are filling faster and the 99 percentile traffic waits longer in the queues even though AQM is trying to maintain the average latency of 25 ms. The gaming traffic in LL SF continues to experience low latency and jitter compared to QB and NQB traffic mix in the CL SF.

**Table 5 LLD ASF Scenario 3 Results**

ASF With BE In LL SF and CL SF									
Scenario	Flow Type	Scheduling Type	Mean Latency (ms)	95 Percentile Latency (ms)	99 Percentile Latency (ms)	Jitter (ms)	US Throughput (Mbps)	DS Throughput (Mbps)	PGS Efficiency (%)
3.00	Low Latency	BE	5.53	7.17	7.28	0.88	1.00	35.00	N/A
3.00	Classic	BE	23.73	40.74	46.21	143.82	19.00		N/A
ASF With PGS in LL SF and BE in CL SF									
3.00	Low Latency	PGS	1.23	1.61	1.74	0.10	1.00	35.00	44.00
3.00	Classic	BE	23.49	42.38	47.46	162.47	19.00		N/A

#### 4.2.4. Scenario 4

In Scenario 4, the traffic adds to burstiness by adding a couple of video conferencing sessions along with frequent FTP file upload sessions. Note that the video conferencing data is passing through the CL SF. The Jitter experienced in CL SF is upwards of 200 ms which is detrimental to gamers QoE and also the video conferencing QoE. The queues are filling to latencies of 50 ms and more. The 99 percentile traffic waits longer in the queues even though AQM is trying to maintain the average target latency of 25 ms. The gaming traffic in LL SF continues to experience low latency and jitter compared to QB and NQB traffic mix in the CL SF, which is the same as the results found in the previous three scenarios.

**Table 6 LLD ASF Scenario 4 Results**

ASF With BE In LL SF and CL SF									
Scenario	Flow Type	Scheduling Type	Mean Latency (ms)	95 Percentile Latency (ms)	99 Percentile Latency (ms)	Jitter (ms)	US Throughput (Mbps)	DS Throughput (Mbps)	PGS Efficiency (%)
4.00	Low Latency	BE	5.58	7.02	7.31	0.88	1.00	39.00	N/A
4.00	Classic	BE	24.01	46.90	57.58	231.19	19.00		N/A
ASF With PGS in LL SF and BE in CL SF									
4.00	Low Latency	PGS	1.16	1.66	1.76	0.10	1.00	41.00	0.44
4.00	Classic	BE	24.27	49.07	58.02	236.24	19.00		N/A


#### 4.2.5. Summary Of LLD ASF Experiments

The following graphs summarize the data in the tables for the four LLD ASF experimental scenarios.


Figures 10 to 13 signify the importance of separating QB and NQB traffic. Scenario 1 acts as the baseline and incremental traffic types are added up to Scenario 4. In general, it can be observed that the 99 Percentile latency of the traffic and the Mean latency and Jitter of the traffic in CL SF is >> than LL SF.

An MSO can choose to use Best Effort as the scheduling type in the Low Latency Service Flow and the Classic Service Flow. The applications that require high throughputs and which are not sensitive to latency can be classified into the CL SF, and they will be treated in a fashion similar to today's latency standards. On the other hand, the applications that are sensitive to latency can be classified into the LL SF, and they will be rewarded with latencies that are much better than today's latency standards. Furthermore, the additional enablement of PGS scheduling within the Low Latency Service Flow will


reduce the latency metrics to millisecond values. But the PGS parameters must be carefully optimized so as to not overgrant or undergrant or a learning mechanism can be introduced to optimize the PGS grant efficiency.




**Figure 10 LLD ASF Both SF BE Mean And 99 Percentile Of Gaming Traffic**



**Figure 11 LLD ASF Both SF BE 99 Percentile And Jitter Of Gaming Traffic**



**Figure 12 LLD ASF LL SF=PGS And CL SF=BE Mean And 99 Percentile Of Gaming Traffic**



**Figure 13 LLD ASF LL SF=PGS And CL SF=BE 99 Percentile And Jitter Of Gaming Traffic**

## 5. Conclusions And Future Work

We proved that the concept of separating QB and NQB traffic assists in NQB traffic achieve lower latency than QB traffic, without degrading the performance of the QB traffic. We conducted experiments with the Proactive Grant Service Scheduling that send guaranteed grants at guaranteed intervals in order to meet the demands of traffic in the LL SF. We observed ~1.5 ms latency and extremely low jitter of 0.10 milliseconds for PGS enabled Low Latency SF within an LLD ASF.

The experiments also showed that either of the scheduling types can be used for LLD ASF – Best Effort or Proactive Grant Service. But it is important to note that Proactive Grant Service can reduce the latency further by proactively granting and reducing the traditional REQ-GRANT time by ~72%!

The CMTS and CM software and the MULPI specification are maturing to achieve the LLD goals. At the time of conducting the study, there were many moving parts because of many different features of the LLD architecture – Queue Policing, Weighted Round Robin, AQM algorithms etc.

There is a need for standard tools to measure latency statistics such as mean, 95 Percentile, 99 Percentile latency and Jitter values. LLD provides a Histogram feature that will give a deeper look in to the service flow buffers and queue build-ups. At the time of writing this paper, we did not have the standard toolset.

In the future, we plan to learn more about tuning of Proactive Grant Service parameters for optimized use of available upstream channel capacity and service flow Maximum Sustained Rate.

Since traffic classification into QB and NQB plays an important role, a future challenge lies ahead for the adoption of traffic classification rules in the LL and CL service flow at the application layer.

This real-world experiment with the CMTS and the LLD capable CM and the bursty traffic generators shows that low bit-rate traffic patterns (i.e. gaming traffic) that behave as NQB can achieve low latency with LLD ASF without impacting QB traffic.

LLD Interops at CableLabs are underway to ensure that the CMTS and the CM equipment are ready for deploying the LLD feature set. Once the spec and software are fully developed and tested, the

introduction of LLD to the consumers will be exciting. LLD is one of the key enablers for the future of DOCSIS in the world of 10G.

## Abbreviations

AMSR	Aggregate Maximum Sustained Rate
ASF	Aggregate Service Flow
BE	Best Effort Scheduling
bps	bits per second
CL	Classic SF
CM	Cable Modem
CMTS	Cable Modem Termination System
DSCP	Differentiated Services Code Point
ECN	Explicit Congestion Notification
FEC	forward error correction
HFC	hybrid fiber-coax
Hz	hertz
ISBE	International Society of Broadband Experts
LLD	Low Latency DOCSIS
LL	Low Latency SF
MSR	Maximum Sustained Rate
NQB	Non-Queue Building traffic
PGS	Proactive Grant Service Scheduling
QP	Queue Protection
QoE	Quality of Experience
QB	Queue Building traffic
SCTE	Society of Cable Telecommunications Engineers
SF	Service Flow
SLA	Service Level Agreement
WRR	Weighted Round Robin

## Bibliography & References

- [1] CableLabs, Low Latency DOCSIS: Technology Overview, G. White, K. Sundaresan, B. Briscoe
- [2] MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I20-200407, April 04, 2020, Cable Television Laboratories, Inc.

# **Enterprise Opportunities Beyond The Pipe**

## **The Second Network**

A Technical Paper prepared for SCTE•ISBE by

**Jay Bestermann**

Sr. Director Managed Networks  
CommScope  
3871 Lakefield Dr Suwanee, GA 30024  
678-473-8153  
Jay.bestermann@commscope.com

**Ken Florenz**

Director of Product Management for ICT and Managed Services  
Altice USA  
1111 Stewart Ave. Bethpage, NY 11714  
516-803-6564  
Kenneth.Florenz@AlticeUSA.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
1. Introduction.....	3
2. What is the second network? .....	4
3. How large is the second network opportunity?.....	5
4. What is required to operationalize the second network?.....	5
5. Altice Business Smart Wi-Fi Case Study .....	8
5.1. Benefits for Altice Business Customer .....	9
5.2. Benefits for Service Providers that Employ Managed Services.....	9
6. Conclusion .....	10
Abbreviations.....	12
Bibliography & References .....	12

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – The First and Second Network.....	4
Figure 2 – Operational Capabilities .....	6
Figure 3 – Managed Services Growth.....	10

# 1. Introduction

For many service operators—service, value, and revenue begin and end with the broadband pipe which is the primary commercial connection between the service provider and the consumer. But as networks evolve and the needs of business customers continue to grow, one of the largest opportunities for increased monthly recurring revenue, churn reduction, and customer service takes off where the first network ends: beyond the pipe.

Service providers are delivering on the promise of the “first network”, broadband, at very large scale, but aren’t fully leveraging their opportunity to deliver the “second network” to their customers. This second network is beyond the pipe and typically consists of router / firewall, switching and Wi-Fi access points to facilitate the day to day operations of the business. In the business segment we estimate that broadband penetration is over 80%, but second network penetration is very low, estimated at less than 10%.

The second network opportunity is significant. As an example, in the SMB market, with less than four access points, there is a willingness to invest by the business. A typical entry broadband circuit for a business will be in the range of \$100 per month. A managed networking offer that consists of a router, switch and APs significantly increases the target addressable market. Business owners are often willing to invest another \$100-\$200 per month in their internal network solution that offers security, wired and wireless connectivity that is proactively monitored and maintained. The monitoring and maintenance of the second network is now business critical to business owners. As bandwidths available to the business regularly exceed 100Mbps, business owners are relying more and more on secure network connectivity to cloud based customer management, accounting, payroll and Point of Sale services for their daily business needs. These cloud-based applications simplify the business owners’ operational requirements and allow them to leverage cloud-based economics for their business applications, but connectivity is no longer a convenience, but a necessity.

Managed enterprise networks involve more than just installing Wi-Fi equipment and incidental hardware for customers. They are in fact fully managed services where the operator is responsible not only for providing and installing the equipment, but customizing it to match the customer's needs, managing it on their behalf, and providing value-added services on top of it. The typical Managed Networks digital experience that provides a sticky engagement with the customer both from a branding point of view as well as useful customer engagement around second network status and analytics specific to their business. With these deeper second network solutions, Service providers can significantly increase their value to their customers across all enterprise business types, large and small. Service providers do have an opportunity to further enhance their value beyond the “Second Network” once they have reached beyond the network edge. This opportunity is delving further into the business with additional value-added services such as video security and other managed IOT offerings that a typical business owner needs but doesn’t have the IT resources to deploy and maintain.

The immediate benefits for operators as well as their business customers are obvious: Customers have access to professionals who manage the entire install, setup and operation of their entire network. Because they don't have to be their own network support organization; they can focus on their core business. A single provider that delivers both the first network and the second give the enterprise a single point of contact if they need assistance with the network or the service. For operators, this value-added relationship not only provides more sticky touch points with the customer, but it also represents the opportunity for revenue through equipment sales as well as recurring revenue through network management. As business networks evolve, so too can operators' role beyond the first network, also known as the pipe.

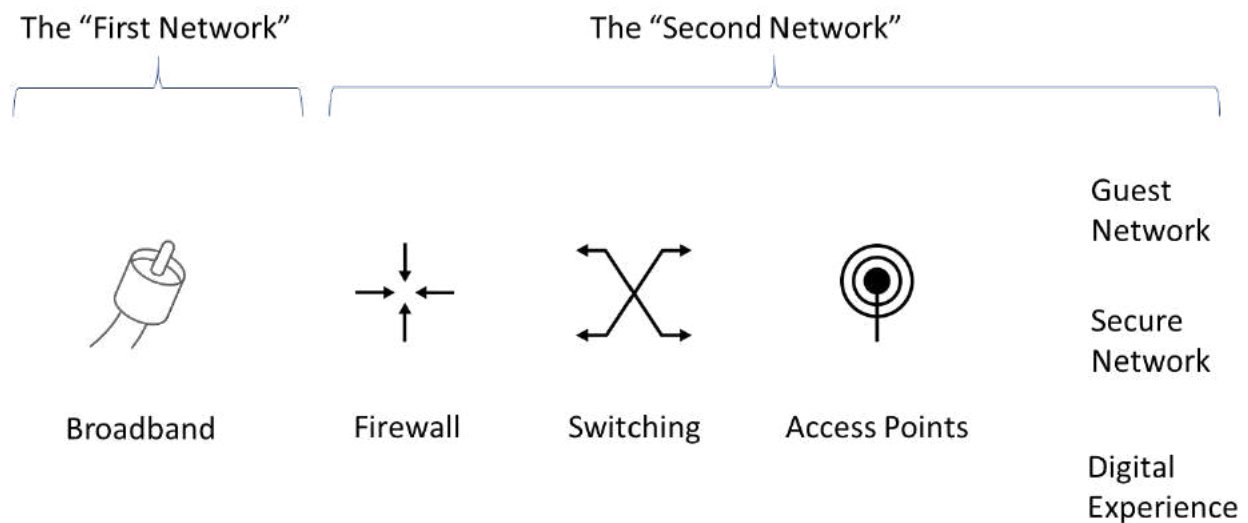


Value added services such as the second network are instrumental in maintaining service provider wallet share. Legacy services including broadband, voice and video are at high and consistent penetration levels across the service provider customer set, which leads to commoditization. Service providers are in a fantastic position to expand their portfolio and increase their customer value. Increased customer value is key to reducing churn and preserving service pricing.

This paper will provide an in-depth review of the market opportunity operational considerations for delivering effective managed enterprise networks. It will also feature a case study with a major US operator to illustrate the challenges and upside in moving past the demarcation point to owning the second network.

## 2. What is the second network?

The second network, depicted in Figure 1 below, is what business users connect their devices to in order to access the internet. This behind the broadband pipe enterprise network is typically composed of a router and firewall platform, Power over Ethernet switching and Wi-Fi access points with enough scale to provide adequate performance and connectivity through the business. This so called second network is viewed by many non-technical users as simply “Wi-Fi”, but people in the know realize that behind every great Wi-Fi experience there’s a fantastic wired backbone. The fact of the matter is every first network, the broadband pipe, is completely unusable without the second network and many times do-it-yourself solutions are under specified, poorly installed and many times not secure. The other fact we all know is the perception of the first network can be severely impacted by the stability and performance of what is deployed behind the pipe. These can be impacted by poor design practices, poor installation practices, inadequate equipment and/or cybersecurity breaches.



**Figure 1 – The First and Second Network**

### **3. How large is the second network opportunity?**

The second network opportunity is very large especially when all applicable vertical markets are considered. The addressable vertical market with a second network offering spans across several different enterprise business types including SMB, hospitality, Multi-dwelling unit, mid-market and large enterprise. For the purpose of this paper we focus on the SMB market in the interest of brevity, but the second network TAM is significant across all verticals.

The target addressable market (TAM) of the small to medium business (SMB) segment alone is approximately \$6B per year for broadband alone. Research group Gartner defines SMB “by the number of employees and annual revenue they have. The attribute used most often is number of employees; small businesses are usually defined as organizations with fewer than 100 employees. The second most popular attribute used to define the SMB market is annual revenue: small business is usually defined as organizations with less than \$50 million in annual revenue.”[1] According to the Small Business Administration there are 30.7 million small businesses in the United States alone employing millions of Americans. Of those small businesses there are approximately 5-6 million that employ people and likely need internet access at their place of business [2]. This segment broadband TAM is in the range of \$500M per month assuming a conservative business circuit price of \$100 per month and 5 million SMB locations that need internet access.

In this SMB segment, the broadband TAM is just scratching the surface of the total networking opportunity. This is because every business needs a highly reliable network on which to run their operation and are more than willing to purchase the solution from their broadband service provider. Our world’s dependence on connectivity has reached utility level especially with the prevalence of cloud-based services for point of sale, customer relationship management and financial accounting. A rough estimate of the second network TAM is estimated to double the broadband TAM at a minimum, perhaps even 1.5 to 2 times the TAM. A rough dollar estimate is \$500M to \$1B per month depending on the product offer or as stated previously \$6B annually. This estimate is limited to the United States due to market data availability, but the true opportunity is worldwide.

Keep in mind that SMB is just one vertical segment that can benefit from a second network offering. The service provider TAM outside of SMB can be expanded to Mid-market, Hospitality, Multi-Dwelling unit, etc.

### **4. What is required to operationalize the second network?**

Providing the second network has many operational considerations that must be addressed as part of a new product launch. Service providers have nationwide scalable expertise in the deployment of circuit related infrastructure, but because most are not broadly addressing the second network they need to thoughtfully prepare. Addressing the second network requires more advanced networking and wireless pre-sales, design, deployment and operational capabilities. The figure below depicts some of the key capabilities that are needed as service providers take their second network journey starting with product planning and moving into product launch and product operations.

Product Definition			Product Launch		Product Operations and Support		
Planning		Design		Deployment	Operations	Support	
Market Analysis		Solution Design		Technical Interview	Core Platform	Tier 1 End User Support	
Technology Evaluation		Technology Selection		Program Management	New Release Mgmt	NOC (Tier 2)	
Product Definition		Solution Integration		Wired Wireless Design	On-Boarding	Tech Support - Product (Tier 3)	
Proof of Concept		Playbook & Process		Equipment Procurement	MACD	Engineering Support (Tier 4)	
Existing Product Audit		Network Security		Product Logistics	Hosting	Repair and Return	
Proof of Concept		Market Trial		Product Installation	Vendor Management	On Site Support	
		Sales Training and Tools		QC & As-built	Reporting and Automation		

**Figure 2 – Operational Capabilities**

As service providers start to address the second network opportunity across verticals, the most careful consideration must be given to defining a very concise product definition tailored to the vertical market as necessary. Each vertical market has different nuances and desires. For example, SMB users want a very simple, reliable solution that is rapidly deployed by the service provider. Mid-Market, MDU or hospitality have more complex requirements and typically require a different product definition. Service Providers are best positioned to address the high-volume opportunities that can be served by a repeatable product definition. The well-defined product with a repeatable architecture and feature set will lead to a more scalable solution that is easier to deploy and support longer term. Managed enterprise services bring the service provider into a realm of a myriad of vendors and solutions that can quickly overwhelm engineering, deployment and operations teams if the service provider doesn't have a structured approach to addressing new Second Network opportunities. There will be a tendency to address every second network opportunity that presents itself, but there must be a method of controlling what non-standard or Individual Case Bases (ICB) projects are pursued. Service providers who establish an ICB control board can determine which of these opportunities will be scalable and repeatable in the future and make good business sense to productize.

Once a clear product definition is established the service provider's engineering organizations supported by the vendor community need to select and integrate a technology platform that will deliver the capabilities of the product definition. This development process can be iterative and time consuming depending on how many different equipment suppliers are involved, the level of pre-integration already performed and the complexity of the vertical market's requirements. There are two portions of a new product launch that are often overlooked are the OSS / BSS order to cash process and the back-office reporting and customer facing self-service dashboard platforms. The integration of these platforms typically involves service provider IT resources and need to be addressed early and often in the development cycle.

Typically, after, but preferably before, system integration efforts are complete and a market trial has been successfully executed, it's time to start training sales, sales engineering, day two support organizations, design and deployment organizations and preparing for life cycle management. Sales and sales engineering training and tools is where a significant investment needs to be made when launching a second network offering. The typical service provider sales teams are used to a very short sales life cycle

and positioning a broadband product that has a pretty clear value proposition and a very limited set of competitors. The second network is very different and requires some level of retooling including training and sales incentive plan adjustments. There are two typical second network products that sales teams will be asked to sell. One will be very standard and meant to serve end customer locations that are less than approximately 15,000 square feet. This product will be easy to position and will typically be sold via an inside sales channel based on the square footage of the business. The second network product to establishments that are over 15,000 square feet is where the sales and sales engineering teams need to be trained and incented to sell. The larger and more complex second network sites will require pre-sales quote tools and outside sales executive and engineering teams that can help interview customers to understand the requirements of their deployment from an existing infrastructure, features and scale point of view. This technical interview will determine if the property fits into the standard product definition, the scale of the deployment, the complexity of any required construction and the wireless coverage areas. From this technical interview the Sales executive and sales engineer have the information they need to create a quote for the prospective customer. A key to collecting this information is well defined survey tools for the account team to leverage as input to budgetary quote tools.

Assuming the outside sales team is effectively trained to perform a survey of the customer and the physical location, the next critical step in the operational journey is to manage and build the complex second network opportunity funnel. The smaller sights are straightforward to managed, but a key to developing and managing this funnel for more complex sites is a rough order of magnitude quoting tool. The rough order of magnitude quote tool removes sales quote friction by providing a simple self-service mechanism to generate complex second network proposals. Traditional manual quoting mechanisms in the complex second network space are too labor intensive and costly, but most importantly slow. If service providers are slow to turn around proposals to their prospective customers, it's very likely they will lose to local competitive options. With an online rough order of magnitude calculator, the account and sales engineering teams can quickly interview a customer and provide a proposal within minutes. Removing this sales friction will drive funnel growth and result in a viable business.

Now that the service provider has a second network product and a sales organization closing business, it's time to start designing and installing new properties. The simple second network that suffices for a less than 15,000 square foot location is sold and delivered on an inside sales and ticket basis. The installation technicians must be trained on equipment installation best practices, deal with sometimes complex ethernet cabling and the nuances of enterprise Wi-Fi access point deployment to maximize coverage and performance. Another key component of the simple second network deployment is trying to simplify scheduling of the deployment. Where possible service providers should build processes to facilitate a single truck roll to install broadband and the second network assuming a new customer acquisition. This will create a better customer experience and reduce expenses related to deploying the new customer. For larger second network deployments that have a varying number of network components such as routers, switches and Wi-Fi Access Points, a design and deployment program management team to assist sales and sales engineering with creation final bill of materials with the assistance of design engineering and ensuring resources are available for the installation are critical. The design and deployment team help shepherd a project from customer contract, through the wireless, physical and logical network design process, the deployment process all the way to ensuring final as built are captured and customer acceptance is received.

As new properties are onboarded, engineering and operations teams that can support past the edge and into the premises with complex Wi-Fi and networking systems that have a variety use cases helps ensure the ongoing customer experience is a high quality one. Tier 1 end user support and network operations teams will need to be trained to now support issues that are past the first network. This is facilitated by as built that are captured by the design and deployment PMO and field services teams during the

installation and quality check process. The as-builts provide network diagrams, network element location, photo evidence of a quality installation, performance results from day of install quality checks, customer acceptance and even a customer survey. This as-built information is an excellent source of information to help Tier 1 and NOC agents perform move, add, change, delete requests, as well as debug and support the second network through its lifecycle. The NOC teams are supporting service delivery and assurance tasks associated with the deployment of properties and have the capability to support end customers that can exist past the circuit. The most critical function of the NOC related to the second network is providing a proactive response to second network outages and anomalies. Traditionally service providers are providing a reactive approach to service assurance, i.e. waiting for the customer to call to report an outage, but a proactively monitored solution is a unique value proposition and a significant differentiator. The proactive monitoring offering is enabled by advanced reporting and analytics that reaches behind the first networking demarcation point and into the second network infrastructure. Advanced second network reporting and analytics at the Tier 1 and NOC level is also critical to implementing a continuous improvement plan around technology and operational execution. This depth of insight is a tremendous value to the second network customer versus the reactive response nature of a typical broadband offering. This enables a one number to call solution for service provider customers, which is a very compelling proposition.

Finally, a retooled cloud operations team is also needed to scale a second network product offering. This team is responsible for operating the back-office technology associated with the second network and advanced debugging of property level issues related to the second network. The Wi-Fi, switching and routing systems ideally have cloud-based systems that provide a second network element management systems, differentiated customer self-service platforms, device authentication, reporting, analytics and other technology associated with the second network. One of the most critical components of the cloud team is facilitating automation of property deployment and flow through provisioning. This is not only critical to scalability, but also repeatability of deployment that leads to lower day two support costs and higher customer satisfaction.

In summary, launching a second network product offering touches nearly every part of the service provider operational organization from sales, to engineering, to operations, to support. This effort is more than worth it given the opportunity to increase the target addressable market and become a more valuable partner beyond the pipe.

## **5. Altice Business Smart Wi-Fi Case Study**

A smart small business is a connected business. Today, both employees and customers not only expect but require reliable connectivity while working, shopping or waiting. A secure private Wi-Fi network now a business-critical item, whether the office footprint is 100 square feet or 100,000 square feet. The service providers are responding to the demand. Small business owners and property managers are focused on obtaining practical, turnkey onsite Wi-Fi capability that just works. Beyond connecting employees securely to boost office productivity, guest Wi-Fi is a customer experience definer that can build positive relationships with visitors and vendors while reinforcing the business brand. In many cases though, it is and can be, the opposite whereby the Wi-Fi is more a detriment due to it being complex and/or underperforming. This will yield the opposite effect business owners are trying to achieve as they are aware the better the Wi-Fi the longer the customer stays, yielding higher revenues.

In addition to serving 4.9 million residences, Altice USA offers voice, data, video and security services to more than 400,000 small and regional commercial businesses, nationwide. Altice Business' managed Smart Wi-Fi solution is a seamless, carrier-managed wired/Wi-Fi network for locations requiring indoor

or outdoor wireless coverage along with wired connectivity, from workspaces to churches, from coffee shops to cultural centers, from medical offices to car dealerships.

Altice Business customers are freed from tedious network configuration tasks, but can offer Guest Wi-Fi to their customers, manage network names (SSIDs) and analyze statistics such as the number of guest Internet users and bandwidth consumed during specific hours of the day and days of the week. The Smart Wi-Fi product enables Altice Business to serve businesses with up to 250 guests and employees and deliver speeds of up to 1 gigabit per second (Gbps), with coverage spanning 24,000 square feet indoors and 12,000 square feet outdoors.

Altice Business Smart Wi-Fi combines both private and guest networks, each with separate access controls. Guest Wi-Fi includes a self-service customizable splash page providing branding and promotional opportunities for the business owner and Altice Business. Content filtering is also offered on the Guest network to enable a safe browsing environment. The business owner can also survey guests as part of the onboarding experience and report on the results of their surveys via the guest experience. For business employees and internal use, Private Wi-Fi includes encrypted private wireless networks for point of sale and other back office use cases in the business. Devices connected to the Private Wi-Fi network can also quickly and safely access wired systems like printers and servers while being protected from cyberattacks from outside the business. Altice Business is helping its customers by creating a high-performance offering that gives SMBs the simplicity and reliability they need to grow their bottom lines without struggling to put people online.

Altice Business has deployed Smart Wi-Fi to thousands of customers across 21 states, performing hundreds of installations per month on the behalf of their customers. This is only achievable due to the simple, effective solution and configuration outlined at the onset of the product development process in conjunction with the various stakeholders.

### **5.1. Benefits for Altice Business Customer**

Business owner's with Altice Business services benefit by having a professional, reliable, and yet simple to use solution whether it be the deployment, integration or usage. It allows the business owner to focus on operating and growing their business and not have the headaches managing the technology as well.


Given the multiple touch point businesses have, it becomes a bit overwhelming at times to identify and connect with the various vendors/providers involved with running a business network. Altice Business has solved for that issue here by provide the necessary technology products and services through one single source.

By offering these services through a single source, it is an enhanced overall customer experience and the provider becomes the trusted advisor to the business.

Customers who now frequent the business recognize the ease and reliability of the Wi-Fi and as mentioned earlier, have a tendency to stay longer periods of time, and purchasing more products/services during that extended timeframe, yielding again a much more enhanced customer experience.

### **5.2. Benefits for Service Providers that Employ Managed Services**

TSIA reports that the Managed Services industry, as evidenced here with Altice USA, is growing by leaps and bounds. Below is a graph outlining managed service growth as a percentage of total revue.



**Figure 3 – Managed Services Growth**

Managed Services with a customer focused outcome through partnerships have yielded a considerable increase in amount of revenue as a result, including ARPU, attach rates, and customer satisfaction. Additional favorable aspects exhibited are reduced churn, minimizing operational expense, increased rate of provisioning by utilizing flow through automation, and the ability to provide additional value add services such as content filtering, additional wireless security features, and reporting capabilities.

Customization for the business owners is a major aspect of the Altice Business' Smart Wi-Fi offering and one that allows the customer's brand to be front and center in the eyes of Altice USA enterprise clientele. This goes to the customer experience aspect of the journey and has assisted in raising Altice USA NPS scores.

While we tout the benefits of all this wonderful collaboration and technology, it doesn't come without some heavy lifting. Training, planning, coordination and systems integration are just a few of the hurdles that teams have overcome. Luckily, the state of technology through the use of APIs and other integration techniques, helps to overcome these challenges.

In the end, the system may be complex, the integration advanced to be able to provide the suite of products and services, but the business, their customers, and our sales teams all view this as a simple, easy to sell, easy to deploy, reliable solution that enables all involved to stay connected, stay informed, and stay focused on their own tasks at hand.

## **6. Conclusion**

Small and Large businesses alike are demanding more from their service provider partners regarding networking solutions. Service providers that offer the second network as part of their portfolio have an opportunity to provide a valuable service to their customer, increase revenue and prevent damage to their brand out of their control caused by substandard networking solutions deployed behind the first network. Second network managed service solutions enable service providers to offer an end-to-end enhanced network experience to their business customers that is second to none. Furthermore, coupling a reliable

second network offering and very highspeed first network positions the service provider for additional value-added services further increasing their customer value and target addressable market.

Service providers that extend their reach beyond the pipe and into the second network will be rewarded with a loyal customer base that is willing to invest in the value they are receiving from their service provider partner.



## Abbreviations

AP	Wi-Fi access point
API	Application Program Interface
ARPU	Average Revenue Per Unit
MDU	Multi-Dwelling unit
NPS	Net Promoter Score
PMO	Program Management Office
SMB	Small to medium Business
TAM	Target addressable Market

## Bibliography & References

[1] <https://www.gartner.com/en/information-technology/glossary/smbs-small-and-midsize-businesses>

[2] <https://cdn.advocacy.sba.gov/wp-content/uploads/2019/04/23142719/2019-Small-Business-Profiles-US.pdf>

# Key Learnings from Comcast's Use of Open Source Software in the Access Network

A Technical Paper prepared for SCTE•ISBE by

**Louis Donofrio**

Sr. Director, Product Management  
Comcast Cable  
1800 Arch Street, Philadelphia, PA 19103  
1-215-283-5023  
louis\_donofrio@comcast.com

**Vignesh Ramamurthy**

Principal Architect  
Infosys Consulting  
1800 Arch Street, Philadelphia, PA 19103  
1-510-529-8155  
vignesh.vr@gmail.com

**Qin Zang**

Software Engineer  
Comcast Cable  
1800 Arch Street, Philadelphia, PA 19103  
1-215-283-6878  
qin\_zang@comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
1. Introduction .....	3
2. Brief History of Open Source Software at Comcast.....	3
3. Open Source Software within NGAN.....	4
4. Telemetry .....	4
5. Containerization .....	6
6. Automation .....	7
7. Network Operating System .....	7
8. Conclusion .....	8
<i>Abbreviations.....</i>	<i>9</i>
<i>Bibliography &amp; References.....</i>	<i>9</i>

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Comcast's Open Source Program Office .....	3
Figure 2 - Grafana dashboard used for network monitoring .....	5
Figure 3 - Original Telemetry Architecture in Comcast DAA .....	5
Figure 4 - New Telemetry Architecture .....	6
Figure 5 - PPOD showing Active and Standby nodes .....	7
Figure 6 - Grafana Dashboard showing ONOS at scale.....	8

## 1. Introduction

This paper will provide insights and lessons-learned from Comcast's Next Generation Access Network (NGAN) team as we have increasingly deployed open source software (OSS) to virtualize the access network. Using GNU/Linux, Kubernetes, Prometheus, OpenFlow and many other open-source software distributions for automation (Ansible), test (Wireshark), troubleshooting (smartmon), and Continuous Integration/ Continuous Development (Jenkins, Git), etc., we have achieved some fantastic results. This did not come without a great deal of heavy lifting, new process development, and flexibility. On occasion, it also required re-direction.

By sharing our experience, we hope that other Multiple System Operators (MSOs) will realize the same cost benefits and new feature speed to market we have seen. By providing details on our software selections, others should be able to achieve an even faster pace of innovation as the cable industry moves towards Network Function Virtualization and Software Defined Networking. Key insights we will share include the detailed network visibility gained by using real-time network monitoring tools like Grafana and Kibana, as well as, requirements to use white box switching equipment, commercial off-the-shelf servers and open source operating systems in the network. Last but not least, we will discuss the importance of actively contributing to the open source community while monitoring community developments and building your own internal expertise.

## 2. Brief History of Open Source Software at Comcast

With the launch of the X1 platform in 2012, Comcast quickly transformed from a multiple system operator dependent on suppliers for innovation to become a software and technology company. Since that time, we have significantly consumed and produced open source software (OSS). Notable Comcast contributions include developing the Reference Design Kit (RDK) platform, which now powers over 60 million consumer premises equipment worldwide. We have also made significant contributions to OpenStack, Apache Traffic Control, and the Apache Hypertext Transfer Protocol (HTTP) Server. OSS users within Comcast include our Video Internet Protocol and Research (VIPER) team who have been using Kubernetes, for cloud DVR and IP Video and have been sharing their insights at conferences including KubeCon and CloudNativeCon since 2016.

In 2016, recognizing the growing importance of OSS, Comcast launched an Open Source Practice Office (OSPO) (figure 1) to drive a more coordinated and planned strategy towards OSS. OSS is a key driver of differentiation, cost reduction and innovation for many leading companies and has become a business imperative for us, as well.



**Figure 1 - Comcast's Open Source Program Office**

Comcast's OSPO's goals include acceleration of software delivery and focus on core differentiation. Comcast's focus is to leverage open standards for innovation and time-to-market while reducing work on commodity non-differentiating software. We are also using OSS to improve cost efficiency and reduce operating expenses while improving software quality, reliability and security.

### 3. Open Source Software within NGAN

Comcast's Next Generation Access Networks (NGAN) team uses open source software (OSS) throughout its Distributed Access architecture (DAA). From Ansible to Wireshark, we are leveraging the innovation and speed-to-market for new feature development that comes from participating in community development.

Four key categories of OSS for NGAN include telemetry, containerized applications, automation and operating systems.

Telemetry tools include Grafana, Kibana, Fluentd and Thanos. Containerization tools include Kubernetes to manage Docker containers which house our virtual Cable Modem Termination (vCMTS) application. For automating runbooks, we have deployed Ansible with zero issues. Finally, our first generation DAA utilizes Ubuntu Linux on our servers and the Open Networking Foundation's Open Network Operating System (ONOS) to manage the control and data plane of our Spine/Leaf switching network.

### 4. Telemetry

Our team realized from the start of the DAA program that the traditional ways of collecting operational statistics regarding the health of our network would not fit the needs of our new architecture. A distributed architecture such as ours was going to be very difficult to troubleshoot and Standard Network Monitoring Protocol's (SNMP's) pull model wouldn't scale for our needs. We wanted to continuously push data to assess the health of the network.

We also wanted to avoid using Command Line Interface (CLI) because they can only be used to check one parameter at a time. With CLI, you need to be very precise with your instructions to get the information needed and there are no graphical views of the information. There were also performance concerns about having too many people logging into critical infrastructure, as well as, security concerns with the likely sharing of passwords.

Using OSS platforms like Prometheus, Grafana, Kibana, and Fluentd for data collection, we were able to provide our operations teams and field personnel with a single pane of glass that provides a view of what is happening across the entire network. Our monitoring probe is recording data every 15 seconds to capture key events and is used to create alarms in Slack to notify support personnel regarding potential issues. Alarms on key parameters are typically triggered after waiting a pre-programmed 5 minutes. As we learn about new types of issues, we continually create alarms to proactively address issues before they become outages.

Any telemetry data being collected can easily be shown in a new Grafana dashboard (figure 2) created as required. Grafana's on-line tutorials make it simple for anyone in NGAN (without programming knowledge or Grafana expertise) to create their own dashboard view.

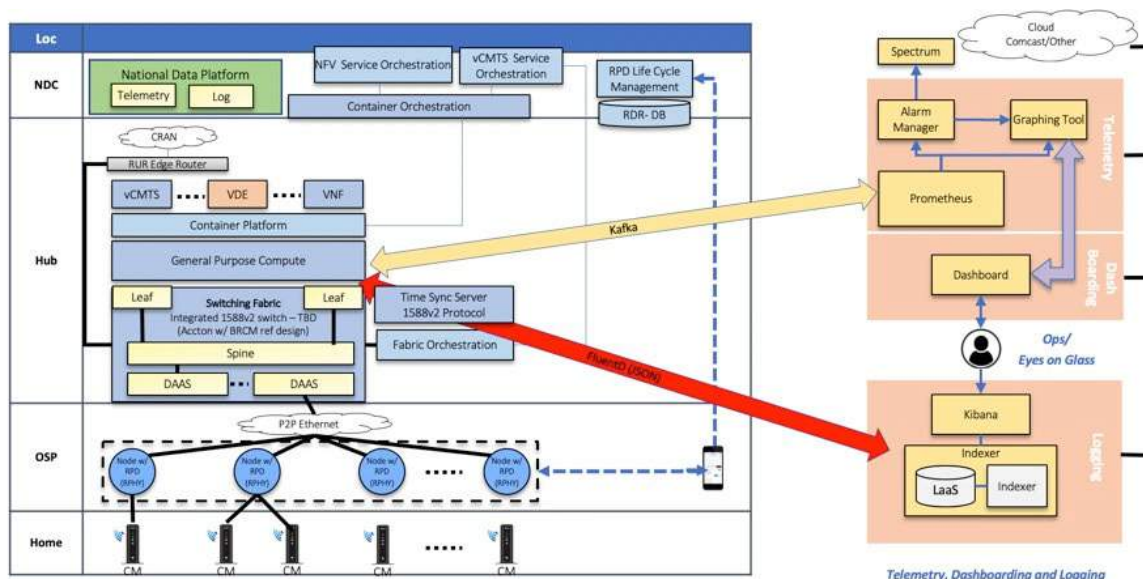
Internally, we have since settled on seven key dashboards to make sense of the DAA environment and they are typically monitored by hundreds of simultaneous users from divisional business leaders to field technicians. Visibility into network performance and understanding of key network parameters has increased dramatically with these tools.



**Figure 2 - Grafana dashboard used for network monitoring**

The increased visibility into network performance has been eye opening. We're now aware of every minor glitch in the network. If a technician opens an optical node in the field, we see it. If power goes out in a particular geography, we are aware ahead of any announcement on the power company's website. One particular use case is that of a supervisor viewing Grafana dashboards in the field to ensure that Remote PHY (RPHY) node cuts are going, as planned.

This architecture was not without its issues. There were gaps in telemetry due to bugs and latency. The initial design (see figure 3) had too many components including stream processors. Stream processors were introduced because the virtual cable modem termination system (vCMTS) was not originally sending data in a Prometheus required format.

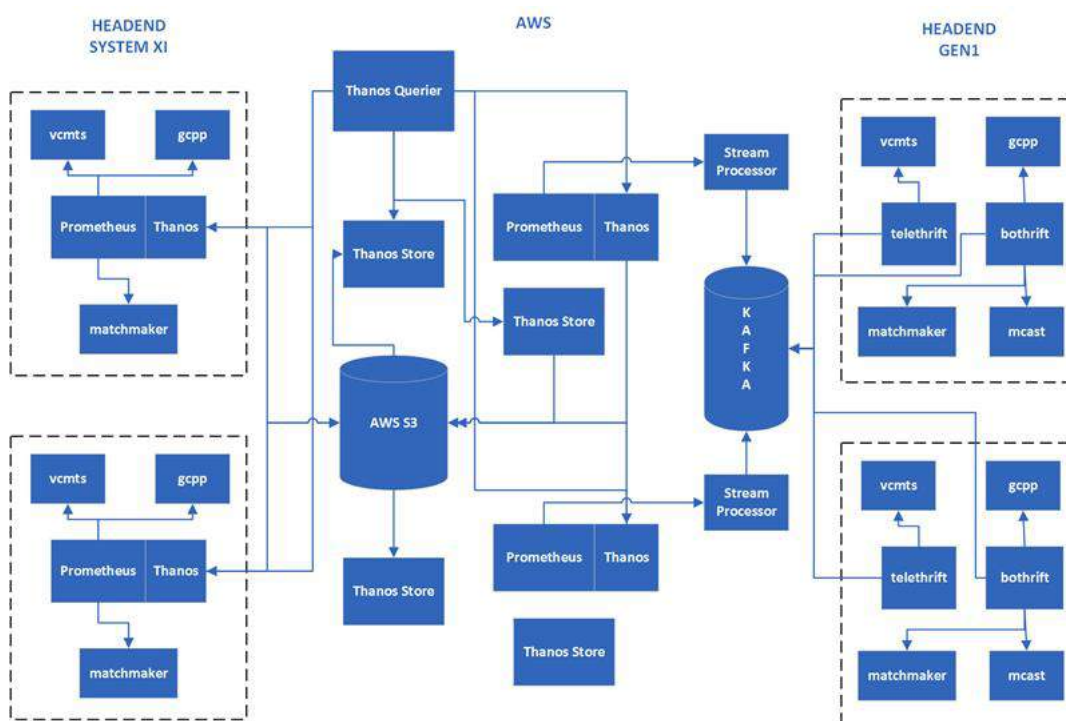


**Figure 3 - Original Telemetry Architecture in Comcast DAA**

Most of our users are looking at Grafana for telemetry. This simultaneous usage does not slow the vCMTS system as telemetry data is coming from stream processors in AWS but the number of users can affect the load of Prometheus which in turn may affect the response time to Grafana. This is something we are addressing in our next generation telemetry solution.

There are storage costs associated with all these metrics and that has to be considered when architecting the solution. We continue to optimize costs for Amazon Web Services (AWS) including limiting metrics storage to 15 days.

Changes were made to improve scalability as we continually add more R-PHY nodes to the network (figure 4).



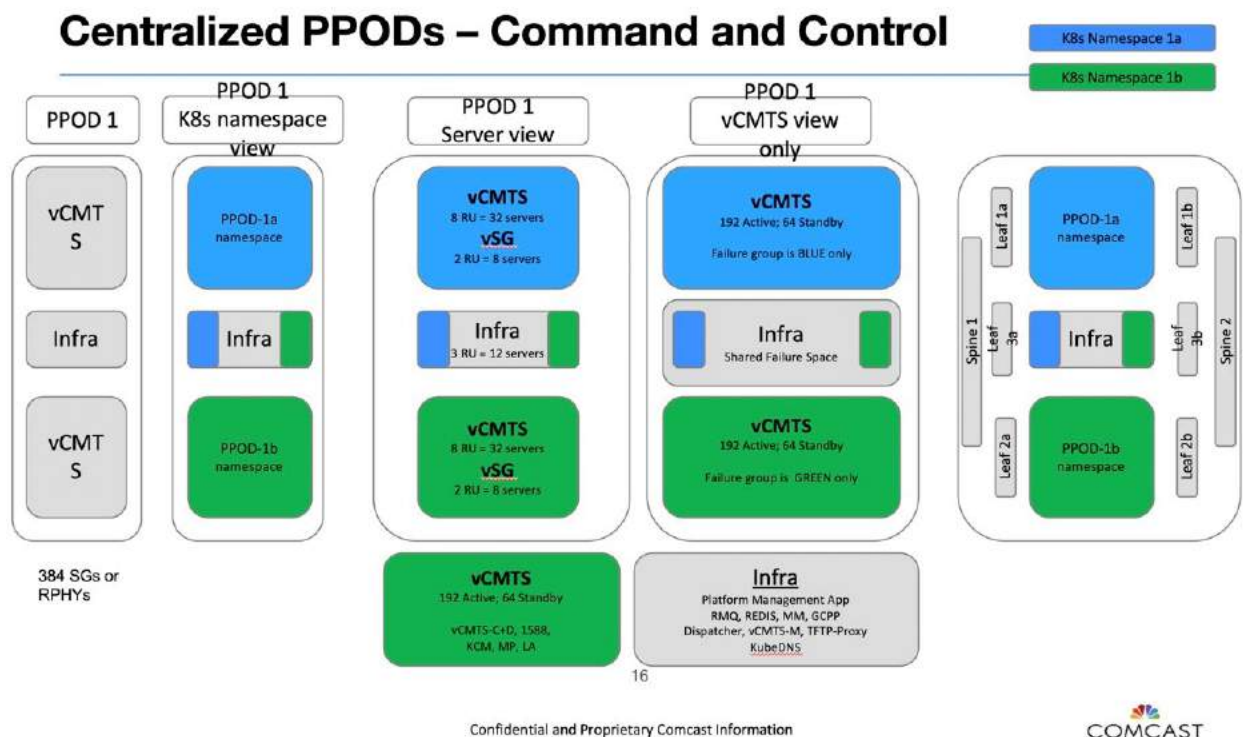
**Figure 4 - New Telemetry Architecture**

## 5. Containerization

Our containers run in Docker using Kubernetes nodes. We initially saw high CPU loads and high input/output queue latency which often resulted in server cartridges hosting vCMTS worker nodes going into read-only mode. This resulted in having to cordon the read-only host, adding additional hardware (figure 5) to increase standby capacity and then fixing the read-only host during a maintenance window. Working closely with our server provider, HPE, we decreased latency of some Input/ Output requests that measured 25 milliseconds to around 0.8 milliseconds by changing Input/ Output timers in Kubernetes, changing queue sizes, and enabling multi-queue devices in Ubuntu Linux.

We also resolved Etcd issues that affected performance. Etcd is often referred to as the brains of the Kubernetes cluster. It is a distributed, reliable key-value store for the most critical data of a distributed system.

Lessons learned include regularly updating the Linux kernels and hardware drivers to take advantage of issues previously solved by the open source community. There has to be a balance between uptime and system maintenance. Also, it is important that your hardware provider understands how you are using their server (Kubernetes, Docker, Ubuntu, etc.) to optimize performance and assist you with choosing the best CPU and SSD for your application and alerting you regarding issues seen by other customers.



**Figure 5 - PPOD showing Active and Standby nodes**

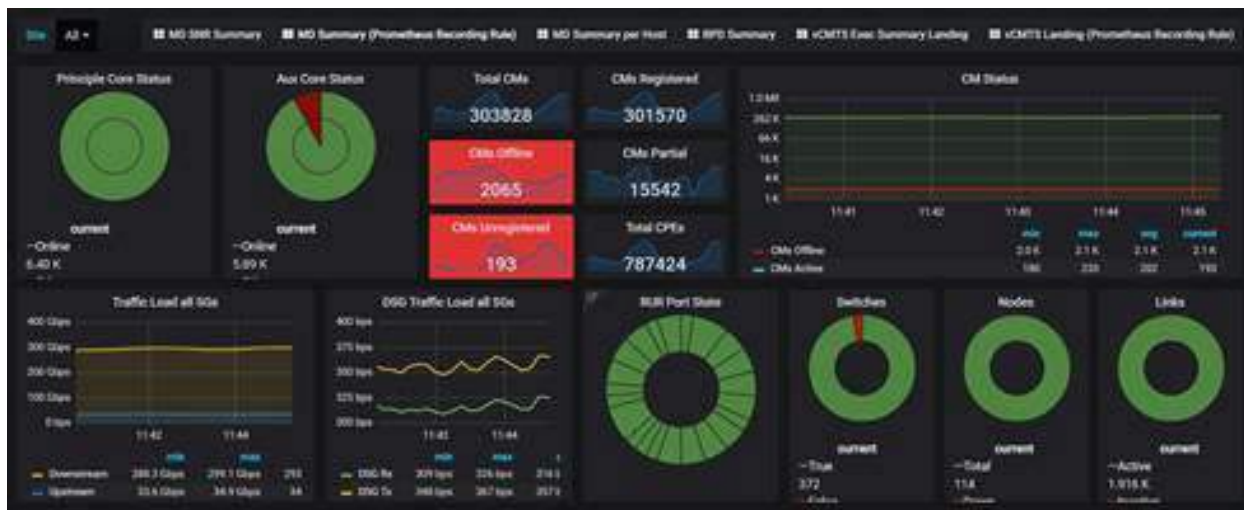
## 6. Automation

We have had nothing but a great experience using Ansible for automation playbooks. Prior to Ansible, our team used hand-crafted scripts which were subject to human error. Other groups within Comcast use Jenkins and Puppet with equal performance.

## 7. Network Operating System

NGAN began using the Open Networking Foundation's (ONF's) Open Network Operating System (ONOS) in production during the first phase of DAA launched in May 2018. With this first generation, we realized significant cost savings in hardware while eliminating software licensing and software assurance costs. Using commercial off-the-shelf servers and white-box switches and working closely with ONF, we achieved limited scale (figure 6) by resolving a number of technical issues.





**Figure 6 - Grafana Dashboard showing ONOS at scale**

Progress included improving reliability during mastership changes, improving multicast handling performance, improving Atomix memory mapping and eliminating silent switch disconnects due to Open Flow socket closures. We achieved up to 40,000 cable modem scale in production by refactoring the route service.

While we had come a long way in terms of stability and performance improvements, ONOS clustering issues still persisted making it unlikely we could achieve our 99.999% uptime goal at 80,000 cable modems scale in time to support our deployment schedules.

Troubleshooting and redundant switch recovery at scale needed additional improvement. To operate at scale, additional features like headless operating mode, a more robust in-band management, handling internal queues, monitoring internal stores, and debugging issues at store level were also needed.

Despite the progress made and the significant investment in gaining ONOS expertise within NGAN, we made the difficult business decision in 2020 to move to a new, generation 2 architecture abandoning the open source network controller in favor of commercial network operating system (NOS)-based switches and a hybrid SDN Controller.

I believe we would have achieved our scaling and reliability goals, but we ran out of time to continue to optimize. We learned a lot in the process and contributed greatly to ONOS community with a number of key innovations. As a result of the progress that Comcast and its partners made commercializing ONOS, it is being evaluated by a number of wireless carriers in the 5G space.

## 8. Conclusion

Comcast was able to utilize OSS to build the access network with vCMTS and R-PHY. We continue to optimize the solution evaluating new hardware and software to improve reliability and scalability at reduced cost per subscriber. Telemetry, containerization and network controller changes were made to improve performance. A switch to a traditional Network operating system was made to immediately reach

scalability and up-time goals. We continue to evaluate new software (including OSS controllers) and hardware technologies to reliably improve cost per bit.

## Abbreviations

5G	5 <sup>th</sup> generation mobile network
AWS	Amazon Web Services
CLI	command line interface
COTS	commercial off-the-shelf (hardware)
DAA	Distributed Access Architecture
GNU	GNU's not Linux
HTTP	Hypertext Transfer Protocol
MSO	multiple system operator
NFV	network function virtualization
NGAN	Next Generation Access Networks
NOS	Network Operating System
ONF	Open Networking Foundation
ONOS	Open Network Operating System
OSPO	(Comcast's) Open Source Program Office
OSS	open source software
PTP	precision timing protocol
R-PHY	remote physical RF layer
SCTE	Society of Cable Telecommunications Engineers
SDN	software defined networking
SNMP	simple network management protocol
SSD	solid state device
vCMTS	virtualized cable modem termination system

## Bibliography & References

<https://www.fiercevideo.com/cable/comcast-backed-rdk-software-reaches-over-60m-devices>

<https://itnext.io/prometheus-for-beginners-5f20c2e89b6c>

<https://wiki.ubuntu.com/Kernel/Reference/IOSchedulers>

<https://medium.com/better-programming/a-closer-look-at-etcd-the-brain-of-a-kubernetes-cluster-788c8ea759a5>

<http://ospo.opensource.comcast.net/about/>

Computer Ganga, “Advantages and disadvantages of command line interface”, 2018. [Online]. Available:

<https://www.computerganga.com/2018/11/Command-line-interface.html> [Accessed 6 August 2020]

# **A Virtual Broadband Network Gateway (vBNG) Approach for Cable Operators in a Distributed Access Environment**

A Technical Paper prepared for SCTE•ISBE by

**Jason Combs**  
Principal Architect  
Comcast  
1800 Arch St Philadelphia, PA 19103  
609-706-9190  
[jason\\_combs@cable.comcast.com](mailto:jason_combs@cable.comcast.com)

# Table of Contents

<b>Title</b>	<b>Page Number</b>
1. Introduction.....	3
2. The challenges of today's access network .....	3
3. A method of network abstraction.....	5
3.1. Service Activation Layer.....	6
3.2. The virtual Broadband Network Gateway structure .....	8
4. Advantages of an NFV and SDN approach compared to a traditional BNG .....	9
5. Conclusion.....	10
Abbreviations .....	10
Bibliography & References.....	11

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 – Evolution of the Access Network .....	4
Figure 2 – Functional components of this network abstraction.....	6
Figure 3 – Interfaces and functions of the Service Activation Layer.....	7
Figure 4 – vBNG structure (OLT example) .....	8

## List of Tables

<b>Title</b>	<b>Page Number</b>
--------------	--------------------

NO TABLE OF FIGURES ENTRIES FOUND.

## 1. Introduction

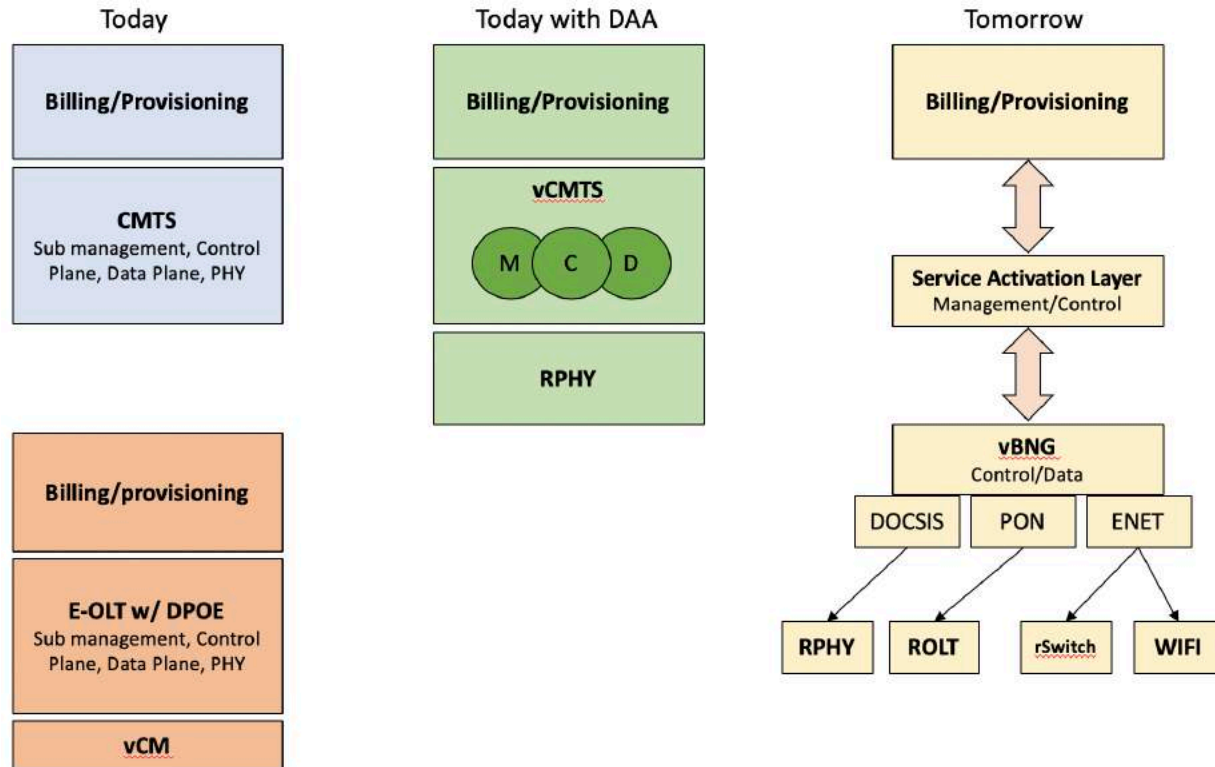
For as long as the Multiple System Operator (MSO) community has built and operated networks, we have been looking for better ways to segment it, and to optimize the work needed to update and maintain it. In addition, serious efforts emerged over the last several years to add other access media to the mix, such as passive optical networks (PONs) and wireless technologies. Among these efforts is network function virtualization (NFV) and software defined networking (SDN), which have captured the attention of the MSOs and of the networking world at large. How does the MSO community take advantage of these concepts to fulfill both today's needs as well as the desire for an easier, quicker deployment of access network technology going forward into the future? In this paper, we will explore:

- Challenges faced by MSOs in deploying new types of access technologies alongside the current (virtual/"v") CMTS
- A method for abstracting service activation to allow for common control of dissimilar access technologies
- A flexible virtual broadband network gateway (vBNG) structure that matches MSO service formats to operate non-DOCSIS access technologies (e.g. PON)
- Advantages of this NFV and SDN approach compared to a traditional hardware BNG

## 2. The challenges of today's access network

The traditional cable operator's access network is centered around the DOCSIS technology that we all know and love. The business support systems (BSS) and the operations support systems (OSS) that are employed are equally centered around supporting the protocols defined by the DOCSIS protocols. CableLabs and we, as a community, have done an excellent job of extending those systems to support our business needs, as our customers and the Internet as a whole evolved from basic data services into voice, IP video, and commercial services.

The focus on a single technology with a single provisioning and operating model led to development of Cable Modem Termination Systems (CMTS) and Optical Line Terminators (OLT) that are fully integrated, as shown in the blue and orange boxes in Figure 1. These systems are typically developed by a single vendor to support all of the service, routing, reporting, and physical functions. The Distributed Access Architecture (DAA) separated the physical generation of the signal, but otherwise remained as integrated as the original CMTS systems. Likewise, the virtual CMTS (vCMTS) has given us a peek into the future by allowing the CMTS functions to live in a server environment, with all of the technology used to accomplish that, but still retains the unified structure seen in prior iterations of the CMTS.



**Figure 1 – Evolution of the Access Network**

The challenge with this integrated structure begins to clearly show itself once we consider providing our services through other access technologies. The expansion into other access technologies is happening for a number of reasons including demand for multi dwelling unit solutions, changes in physical construction costs, and the desire to compete in the Wi-Fi and wireless arenas. Although more distinctly seen now, these issues have been lurking in the shadows all along. The DOCSIS service, provisioning, and operational models do not mesh well when combined with the wireless, fiber, and ethernet technologies that allow us to expand our network in non-traditional directions. On the other side of the coin, the BSS and OSS systems that have grown up in reaction to the DOCSIS architecture force us to attempt to have other access technologies mimic that architecture to avoid major back office changes. The following paragraphs provide a brief synopsis of five admittedly intertwined priorities that are difficult to achieve with the current structure.

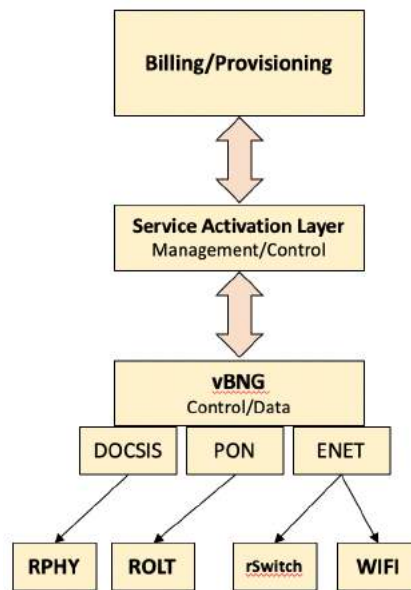
- Swift integration of new access technologies. Beyond the need to develop a technology for our physical network, new technologies typically also come with a full system design and integration. All of these system components require time to develop, and often duplicate the same functions from other systems. Most introductions of new vendors cause us to start this cycle again, representing a high barrier to entry and ultimately fewer options for technology partners.
- Consistent services provisioning. It is cumbersome to develop and maintain separate BSS systems, in whole or in part, for every technology that we deploy. The problem is a compounding one. It is a massive effort to deploy new authentication and service definition systems, and once they exist, they are difficult to maintain while assuring that service definitions remain aligned

with the existing DOCSIS systems. The way around this problem is to create a translation system, such as the virtual cable modem introduced by DOCSIS Provisioning of EPON (DPoE). This, too, needs to be maintained and updated with the latest changes, and can easily fall behind. Workaround remedies add time to any process and represent custom work for the MSO community.

- Service consistency. DOCSIS provides a robust set of service features, which are implemented in its service flows with extensive classifiers, subscriber management filters, and frame accounting with Simple Network Management Protocol (SNMP) and IP detail records (IPDR). Other access technologies have similarly robust features that are simply different from what we typically use. Aligning these methods can cause hardware and feature control issues that limit our choices. The worst-case scenario is the network could look and act differently for one customer than it does for another.
- Testing velocity. Swift testing is an inherent challenge with any additional development. What exacerbates it is the unnecessary replication of the same function within multiple systems. Rather than testing the unique functions, much time is spent retesting different implementations of the same function.
- Operational model differences. With different systems come different interpretations of specifications and different interface models. This is aggravated by specifications created by organizations with very different goals, that apply to other access technologies. They also create additional work, to translate from one system to another.

### 3. A method of network abstraction

There is no simple or fast solution to the challenges stated in the previous section. A new system organizational abstraction is needed to make significant progress in resolving these issues. The existing system needs to be broken into distinct components that can evolve in isolation from each other. In kind, these components need to be linked by extensible application programming interfaces (APIs) that can be modified to meet the needs of the future. With this isolation and extensibility, we remove the need to replicate functions that are common to all of the access technologies and give ourselves the flexibility to integrate any technology quickly. Figure 2 below represents the proposed network abstraction.



**Figure 2 – Functional components of this network abstraction**

This architecture focuses on three major areas of providing services through a scaled access system:

The billing and provisioning systems manage the customer entitlements and interactions with us as cable operators. It is necessary to break away from the tight DOCSIS integration that our back-office systems are built on. Much of the goal of this architecture is to enable that transition. The change itself is a massive amount of work that is not directly addressed within this paper.

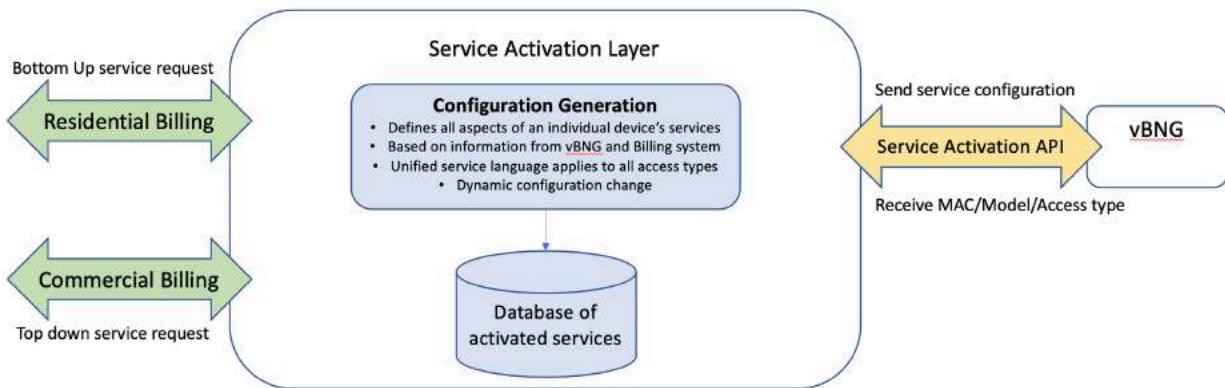
The service activation layer receives the entitlements from the BSS and creates the detailed service description for a particular customer device.

The vBNG implements the service description provided by the service activation layer and manages the moment by moment experience of the customer. This layer can be broken down further, which we will explore in the following sections. We can add a fourth area to this list of three areas to include the operations and data gathering systems. This has been explored by prior papers in detail, for example, in the SCTE paper “The Future of Operations: Building a Data-Driven Strategy” [1] and will not be explored here.

### **3.1. Service Activation Layer**

The service activation layer exists to create a means of defining the customer’s access network services that is abstracted away from both the BSS system and the access technologies that those services ride on. There are several aspects to this component that are necessary to be able to provide this function as illustrated in Figure 3.





**Figure 3 – Interfaces and functions of the Service Activation Layer**

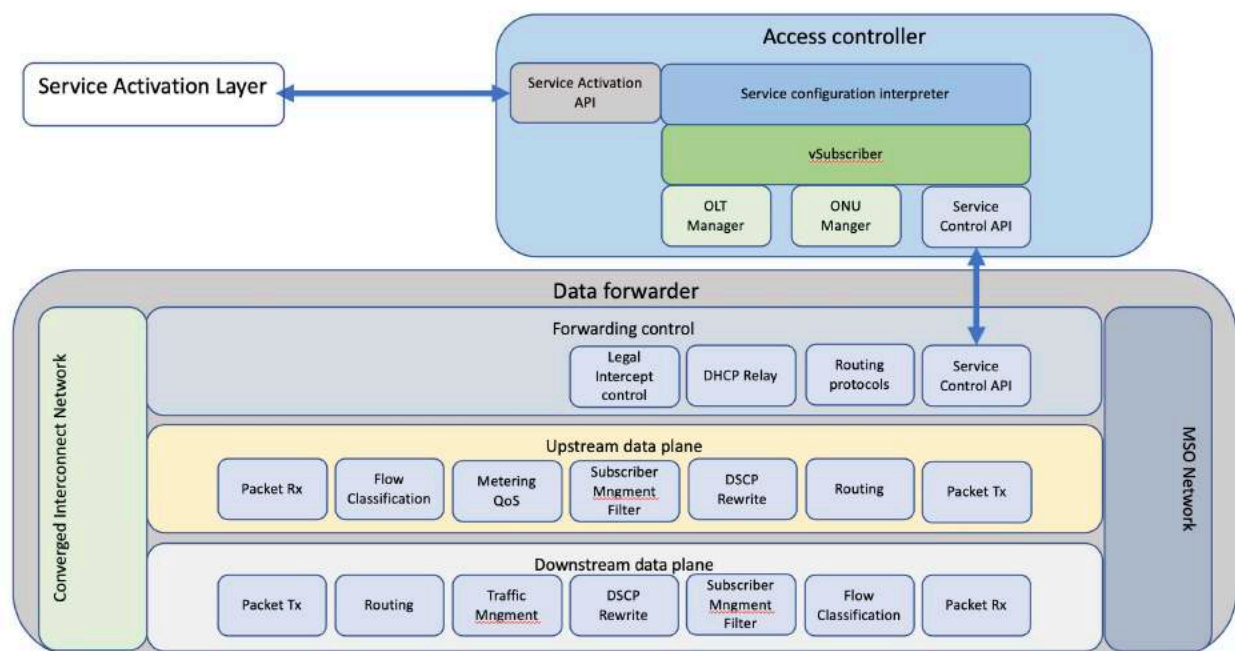
- Service Entitlement flexibility. Many access networks, such as DOCSIS, use a bottom up provisioning methodology where the end device requests service. However, there are examples of top down access networks, such as commercial ethernet and residential ethernet, where the end device does not include this capability. It is important to have an API that allows for both methods of entitlement and configuration.
- Unified Access Service Descriptors. These descriptors are imagined to be in the form of a YANG data model to communicate the services to the vBNG deterministically. They would be formed via a configuration generator that takes the access device type and model along with the customer's entitlements into account. With a common vBNG data plane and a desire to have consistent services, it is imagined that the YANG model for the customer's services would be largely common between access types. The Virtual Provisioning Interfaces Technical Report [2] produced by CableLabs has defined a YANG model that may be a good fit for this language.
- Access/Customer Premise specific configuration modifiers. While the majority of the service definitions will be common, there will certainly be access-specific configurations that would be of interest. For example, a PON system may want to be able to configure optical parameters that do not apply to other access technologies. Access-specific extensions of the universal YANG model would cover this case.
- Database of activated services. With millions of devices in our networks, it is often difficult to have a definitive knowledge of the exact service configuration the customer device is operating upon when a problem is discovered. The service activation layer should record the exact configuration given for each device to aid in troubleshooting and to add traceability.
- Push/Pull of configuration. In the same vein as supporting both bottom up and top down provisioning, this is the mechanism for providing the YANG model configuration to the vBNG. This also enables the ability to change customer device configuration dynamically.

Introducing the service activation layer gives us several capabilities and advantages that help us overcome the problems defined earlier in this paper.

- It abstracts billing and provisioning away from the configuration of the network.
- It presents a consistent service provisioning mechanism for all access networks and all services running through them.
- It brings consistency to service definitions across all access technologies, making their implementation also consistent, within the vBNG
- It eases the application of dynamic services to existing devices and allows for modification of existing services.

### 3.2. The virtual Broadband Network Gateway structure

Many of the issues discussed here can be resolved with a proper vBNG implementation. There are a number of options, in terms of how the vBNG can be laid out, and which components are included. Let us start with a description of the basic components in this method.



**Figure 4 – vBNG structure (OLT example)**

Access controller. This sets up the access network itself and the services that flow through it. To accomplish this task there are several steps that are necessary. A service configuration interpreter translates the YANG model produced by the service activation layer into access-specific configurations used by the vSubscriber data construct. The access manager, displayed as an “OLT” or “ONU Manager” in the diagram, is responsible for configuring all remote access components. A service control API then sends the necessary configuration information to the data forwarder, which is where components common to all access networks are configured.

*Data forwarder.* This provides frame processing and basic subscriber services. Certain functions can be enabled or disabled, such as a traffic manager, for access networks that do not have a built-in means of doing so. On a per customer basis, features like subscriber management filtering and DSCP rewrite can be enabled or disabled based on the services defined via the service activation layer. To aid in effective use of network and processor resources, the data forwarder can be implemented as a single unit or can be separated out for greater flexibility. The first logical separation is the network control, which includes the functions necessary to connect to the MSO network as well as typical services such as DHCP relay, routing protocols, and legal intercept controls. The upstream data plane is the next separation and provides data forwarding services from the customer into the network. Upstream traffic typically has a lower volume expectation. Finally, the downstream data plane provides data forwarding services from the network to the customer. In this case, all access networks would utilize the traffic manager function to shape traffic going downstream. Downstream traffic typically has a higher volume expectation.

These components give us several advantages, no matter how they are laid out in a virtualized system.

- The data forwarder includes common code for all access networks. This is a major facilitator of service consistency, as all customers are served by the same implementation for many of functions provided. From a development perspective this allows us to develop once but use the code many times. Adding new routing protocols, for example, only requires the single implementation to be tested. There is also no need to test the common code extensively when a new access network is developed.
- The Service control API is well defined between the access controller and the data forwarder. This allows for easy and unambiguous integration with new access controllers.
- New development for a new access technology is primarily limited to the access specific functionality such as the access controller, the remote access element, and interoperability with in-home equipment.

## **4. Advantages of an NFV and SDN approach compared to a traditional BNG**

The traditional BNG is the staple of many telecommunication companies' network and has proven to be a solid solution for their access needs. So, why not use the same thing? The answer lies in flexibility. Just as with traditional CMTS deployments, a network operator must choose and size the BNG appropriately for both their current and future needs. Inevitably, when the network capacity or business needs surpass that BNG's capabilities, the operator must upgrade or add more BNGs -- or worse, do both.

Flexibility and reuse are where NFV and SDN shine. This is exemplified by allowing the system to expand functionality naturally into new access networks, while also allowing each of the components to continue to evolve -- to remain the best they can be, without requiring that the entire system be recreated. In terms of capacity, the system relies on generic processing hardware to do its job.

The first advantage with this approach is the ability to directly add capacity by adding more compute power, without touching the rest of the system. Secondly, when using generic hardware, the system can be laid out in different ways to fit the most effective technology, in terms of cost and efficiency.

There are several ways in which these components can be laid out in a virtualized system to optimize the available technology.

The first decision for optimization is whether one instance of the vBNG will handle one service group or many service groups. A single service group vBNG gives you the greatest control and isolation at the expense of more wasted processing cycles, while a multiple service group vBNG gives you more efficient processor utilization at the expense of more complex control and isolation.

The other area of opportunity is whether and how to separate the components of the vBNG to more efficiently assign processing resources. Separating the access controller from the data forwarder is attractive as the access controller and the data forwarder have very different functional and service assurance metrics. This separation would allow for a considerable amount of processor utilization optimization. Perhaps this would give us the ability to have the access controller live higher up in the cloud. Another way to optimize the component separation is to divide the upstream and downstream data planes. This would allow us to take advantage of inherent differences in upstream and downstream utilization. The upstream data plane has significantly less usage, even with symmetric access technologies, and can be multiplexed more effectively, while the downstream data plane has significantly more usage and needs more network and processor resources to be assigned.

Details of these approaches are a subject of considerable discussion in and of themselves. Intel has published an architectural study that dives deep into the factors of these decisions [3] that may be interesting to the reader of this paper.

## 5. Conclusion

The protocols defined in the DOCSIS specifications have led us down a path that has centered our back office and access networks around those protocols. This has been very good to us, but it is time to integrate other access technologies into our portfolios. The NFV and SDN evolution gives us an opportunity to refactor the way that we build and run our networks to support both the network of today and the multi-technology network of tomorrow.

In light of this goal, this paper has described an access network abstraction which defines a clear demarcation of system functions between the BSS and OSS, service activation, and network access. This architecture should allow us to:

- Deploy new technology quickly.
- Provision access networks in a unified way.
- Have consistent services.
- Minimize testing.
- Enjoy one operational model.
- Evolve each component of our network without interfering with the others.

We must not only produce a network that solves the problems that we can foresee, but one that allows us to continue to adapt well into the future.

## Abbreviations

vBNG	virtual broadband network gateway
MSO	multiple system operator
NFV	network function virtualization
SDN	software defined networking
PON	passive optical network

CMTS	cable modem termination system
vCMTS	virtual cable modem termination system
DOCSIS	data over cable service interface specification
BSS	business support system
OSS	operational support system
DAA	distributed access architecture
DPoE	DOCSIS provisioning of EPON
vCM	virtual cable modem
ENET	ethernet
RPHY	remote PHY
ROLT	remote optical line terminator
SNMP	simple network management protocol
IPDR	IP data records
API	application programming interface
YANG	yet another next generation
OLT	optical line terminator
ONU	optical network unit
DSCP	differentiated services code point
rSwitch	Remote Switch or field deployed switch
SCTE	Society of Cable Telecommunications Engineers

## Bibliography & References

- [1] SCTE paper “The Future of Operations: Building a Data-Driven Strategy”
- [2] Virtual Provisioning Interfaces Technical Report - CableLabs
- [3] <https://www.intel.com/content/dam/www/public/us/en/documents/platform-briefs/broadband-network-gateway-architecture-study.pdf>

# The Future of Cable Television Audio is Accessible

A Technical Paper prepared for SCTE•ISBE by

**Mark Francisco**

Fellow

Comcast Cable

1800 Arch Street Philadelphia, PA 19103

(215)286-8959

Mark\_francisco@comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
1. Introduction.....	3
2. What we must do – Regulations for MVPDs and Broadcasters .....	3
3. Conventions in Cable Television Audio .....	4
3.1. Retransmission.....	4
3.2. Conventions in Streaming Media Provider Audio .....	5
3.3. Conventions in Cable Television Broadcast Programming.....	5
4. Where Do We Go From Here.....	7
4.1. The United Nations Already Solved This Problem .....	7
4.2. Accessibility is About More Than Just Content Signaling .....	7
4.3. The Future of Over-the-Air Broadcasting .....	8
4.4. M&E+D.....	8
4.5. Why is this important.....	9
4.6. Automating Audio Description.....	10
5. Beyond Audio Description .....	10
6. Conclusion.....	10
Abbreviations .....	10
Bibliography & References.....	11

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1. Audio Ducking Example.....	8
Figure 2. Multiple Audio Presentations Through AC-4 Substreams .....	9

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 – Streaming Offerings with Audio Descriptions .....	5
Table 2 - Broadcast Offerings with Audio Descriptions .....	6
Table 3 – MPEG-2 Transport Stream Audio Type Values.....	7

# 1. Introduction

Television sound hasn't changed much since its first appeared in the U.S. in 1941. In the 1940s, television audio was broadcast as a frequency modulated (FM) subcarrier to the video program. A stereo subcarrier was added in the 1970s. A secondary audio channel was introduced in the 1980s, primarily for dominant non-native languages such as Spanish. Digital audio introduced surround sound over cable and satellite services in the late 1990s, and over-the-air transmission with the conversion to Advanced Television Systems Committee (ATSC) broadcasting in the 2000 aughts. Twenty years later, despite the explosion of Internet Protocol (IP) and Internet delivery, television audio has made little progress, with the exception of the broadcast of Descriptive Audio or Video Description (formally named as Audio Descriptions) and the expansion of two dimensional audio to three dimensions with object-based audio, the latter of which is only available using IP-delivery and Blu-Ray. The future of broadcast television has been standardized internationally. with the recent release of ATSC 3.0. Its advances include video and audio formats by moving to IP encapsulation, adding high-dynamic range and UltraHD (4K) video, and advancing audio to a novel audio ecosystem, AC-4 [1], supporting greater efficiency, fidelity and personalization compared to current audio delivery.

Since the advent of “talkies,” in the late 1800s, audio has always been an appreciated accompaniment to video. There is a long-standing truism that “sound is more than half the picture”. Despite its ability to convey a story comprehensively (as in radio), however, television sound is seldom delivered on its own.

For 12 million people in the US [2], sound conveys the entire television experience. Sight impairment and blindness affects a large and growing segment of the population. Congenital conditions, disease, injuries and age fuel the reliance on audio over visual inputs, and these causes are predicted to grow over time. Conversely, the same is true of those who rely partly or completely on vision over hearing. The need for information, communications and entertainment in video- or sound-only format is increasing.

The desire to advance single sensory experiences is not just a regulatory or moral imperative -- it is good business. Inclusive products yield business opportunities for disabled and abled people alike. Disability is not necessarily a permanent condition -- it is defined by the World Health Organization [3] as a mismatch between the individual's ability at the moment and the environment with which they are in. Therefore, we may all experience disabilities (and need for accommodation) at some point in our lives.

While much has been accomplished in the pursuit of equivalent experiences for people of all abilities, many opportunities exist. Cable television is a sexagenarian industry with incredible establishment, reach and constraints preventing evolution. This fact can only serve to increase resolve, as disruptive change is possible and can offer large advances in value, equitability and experience.

## 2. What we must do – Regulations for MVPDs and Broadcasters

The Twenty-First Century Communications and Video Accessibility Act of 2010 [4] resulted in the creation of a set of Federal Communications Commission (FCC) regulations. Enacted in 2011, these regulations require commercial television stations that are affiliated with one of the top four commercial television broadcast networks (ABC, CBS, Fox, and NBC) and are located in the top 60 television markets to provide 50 hours of video-described programming per calendar quarter during prime time or on children's programming, as well as an additional 37.5 hours of video-described programming per calendar quarter at any time between 6 a.m. and midnight. In addition, multichannel video programming distributor (MVPD) systems that serve 50,000 or more subscribers must provide 50 hours of video description per calendar quarter during prime time or on children's programming, as well as an additional 37.5 hours of video description per calendar quarter at any time between 6 a.m. and midnight, for each of



the top five national nonbroadcast networks that they carry on those systems. The top five nonbroadcast networks currently subject to the video description requirements are USA Network, HGTV, TBS, Discovery, and History. The FCC published a Notice of Proposed Rulemaking (NPRM) on April 22, 2020 proposing an expansion from the top 60 markets by 10 annually for four years beginning in 2021. [5]

The NPRM also recommends changing the term for video-described content to “audio description” from the current convention of “video description”. Additionally, the apparatus must provide a simple and easy-to-use mechanism for activating the secondary audio stream for audible emergency information.

Manufacturers must provide access to video description through a mechanism that is reasonably comparable to a button, key, or icon.

Digital apparatus must be designed so that control of appropriate built-in functions included in the digital apparatus (i.e., user interfaces) are accessible to and usable by individuals who are blind or visually impaired, if achievable. Specific requirements exist with respect to on-screen text menus and other built-in digital apparatus functions.

Manufacturers of digital apparatus must ensure that information and documentation it provides to its customers is accessible, if possible (e.g., user guides, bills, etc.). Digital apparatus manufacturers also must comply with requirements relating to the manner in which they notify consumers that digital apparatus with the required accessibility features are available.

### **3. Conventions in Cable Television Audio**

Cable television is accessible. As a primary and often lifeline service to its customers, cable, satellite and fiber providers of multi-channel television services are responsive to regulation, social need and the business opportunity that serving the entire market can bring. Being accessible can often result in unanticipated benefits in what is not traditionally considered disability usage. An example is the use of closed captioning in health clubs and bar/restaurants to benefit patrons who are unable to hear due to the high ambient noise level. Cable television companies are subject to regulation by the FCC, state regulators and franchisers and must adhere to content carriage contracts. The result is high availability of accessible content and services. Many do not stop at the minimum bar of regulation, and add additional features that benefit people with disabilities beyond the law. Examples include the Comcast Accessible Remote and the Altice One Voice Remote, among others. The benefits are significant to the populations served, be they with disabilities or without, and to brand identity.

#### **3.1. Retransmission**

Either by convention or contract, broadcast over-the-air networks and cable channels are typically retransmitted in their as-received state, that is MPEG-2 Transport Stream with MPEG-2 video compression and AC-3 audio compression. Approximately 2/3 of cable television stations are received in 1080i (1920 x 1080 interlaced, 30 fps) video, 5.1 channel primary audio with AC-3 compression, monaural secondary (SAP) audio with AC-3 compression and CTA-608 closed captions wrapped in CTA-708 transport. The majority of the remaining 1/3 of stations are 720p (1280x720 progressive, 60 fps) video, 5.1 channel primary audio with AC-3 compression, monaural secondary audio program (SAP) with AC-3 compression and CTA-608 closed captions wrapped in CTA-708 transport. A small number of stations are received as standard definition 480i (640 x 480 interlaced, 30 fps) video, 2.0 channel primary audio with AC-3 compression, monaural secondary (SAP) audio with AC-3 compression and CTA-608 closed captions wrapped in CTA-708 transport. This is almost identical to the ATSC over-the-air

broadcast format, with the exception of the vestigial sideband (VSB) modulation required for radio transmission. Increasingly, MVPDs are converting the received video compression to a more efficient H.264 while retaining the audio bitstreams unchanged. Typical as-delivered cable channels require 8-12 Mbps for MPEG-2 video (less for standard definition), 3-6 Mbps for H.264 video and 384-512 kbps for two audio services.

Due to constraints in receivers and the inability for broadcasters to dynamically change descriptors, the primary and secondary audio streams, when both are present, have differing but no more than two language codes. This results in the typical practices of signaling the primary language as English (ISO\_639\_Language code ENG) and secondary audio program language of Spanish (SPA), Portuguese (POR), French (FRA) or Middle English (ENM). The secondary language codes are often not representative of the audio program's actual language, potentially resulting in viewer confusion. Furthermore, broadcasters that deliver audio descriptions (AD) can change the service provided on the secondary audio program between a secondary language such as Spanish and a complete mix of English program with English Audio Descriptions but retain a consistent but incorrect SAP language code of Spanish.

### 3.2. Conventions in Streaming Media Provider Audio

Streaming media providers are content creator, aggregator, and hosting entities that rely on user-supplied internet network connections. They deliver content using Hypertext Transfer Protocols (HTTP) protocols such as HTTP Live Streaming (HLS) or MPEG-DASH (dynamic adaptive streaming over HTTP). These relatively modern packaging methods are more flexible in terms of variety of content presentations, and this can include multiple audio and video formats. These providers deliver almost exclusively video on demand, rather than broadcast formats, so they use file-based rather than live workflows. This agility and lack of linear complexity results in a much greater offering of audio selections on a per-asset basis. A recent survey of popular streaming media providers discovered thousands of assets containing audio descriptions as an audio option. Many are delivered as a complete mix in surround sound format, rather than monaural. The most advanced offerings are found on Apple TV+ which includes several languages of audio descriptions in object-based audio format (Dolby ATMOS). Table 1 lists the number of titles available with audio descriptions from popular streaming media providers. [6]

**Table 1 – Streaming Offerings with Audio Descriptions**

Provider	Count (with AD)
Apple TV+	32
iTunes	1341
Disney+	589
Hulu	91
Netflix	1271
Amazon Prime	1831

### 3.3. Conventions in Cable Television Broadcast Programming

A check of a single day's (August 5, 2020) programming listed 109 shows with audio descriptions. (Table 2)

**Table 2 - Broadcast Offerings with Audio Descriptions**

<b>Broadcast Network</b>	<b>with AD</b>
ABC	28
CBS	13
Discovery	25
Fox	25
HGTV	35
History	28
NBC	18
Oxygen	4*
SyFy	2*
TBS	19
Telemundo	7*
CW	5*
TNT	12*
USA	20

\* denotes networks without AD mandates

These numbers are substantial, and they are generally increasing. The challenge remains in finding specific programming with AD, and then selecting the proper audio service when desired, and deselecting it when not. Most households with visually impaired residents are shared with household members who lack vision challenges, and these individuals may not want audio descriptions present while viewing programming.

Broadcast television programming audio services are ambiguous for a number of reasons. Due to constraints discussed in section 3.1, cable channels are limited to two audio services and the two must be identified with differing languages. Audio descriptions are conventionally placed in the SAP, often displacing a second language such as Spanish. Broadcasters do not dynamically change language codes, which are part of the MPEG-2 transport stream that is created by the broadcast contribution encoder. Therefore, if a broadcaster occasionally transmits Spanish as SAP and occasionally transmits AD as SAP, the broadcast language (ISO\_639\_Language) is always set to the SPA denoting Spanish. In the U.S. a workaround is often used that signals AD as Portuguese (ISO 639 language code POR), and in Canada, Middle English (ENM) and Middle French (FRM) is used. If a cable service-provided user interface allows a user to select from the languages signaled by the network broadcaster, audio descriptions are almost always incorrectly identified. The network also supplies metadata to populate program guides, and these often have correct information about audio services, but only for time of broadcast. Shows are increasingly viewed in a time-shifted fashion through digital video recording, start-over and on demand methods, and these may differ in provided audio services. Programs that are broadcast with audio descriptions as SAP are not available with SAP when viewed on demand. MVPDs have increasingly been improving their user experiences, using logic to assist with selecting audio descriptions. Oftentimes, program metadata is considered first, followed by selecting SAP when the program data is not affirmative as the existence of AD.

The ambiguities mentioned above do not apply to streaming media providers. There, the content hosting and user experience is tightly bound, and content is accurately and affirmatively identified in the file-based metadata and presentation information.

## 4. Where Do We Go From Here

Technology is quick to change, standards are very slow to evolve, and consumer behaviors range the entirety of the space between. The technologies to improve the user experience around accessibility of content are with us already, and some of them have been defined in standards for some time. Let's look at what can be done now and what we can anticipate in the near future.

### 4.1. The United Nations Already Solved This Problem

The International Standards Organization (ISO) of the United Nations includes the Moving Picture Experts Group (MPEG). The system standard for MPEG, including the basis for ATSC and Digital cable transport streams (ISO13818-1), was released in the late 1990s [7]. Version 2, released in late 2000, contains a second field to the ISO 639 [8] language descriptor entitled **audio\_type**. The values for audio type are located in Table 3.

**Table 3 – MPEG-2 Transport Stream Audio Type Values**

Value	Description
0x00	Undefined
0x01	Clean effects
0x02	Hearing impaired
0x03	Visual impaired commentary
0x04-0xFF	Reserved

Inserting the code 0x03 in the MPEG-2 headers will result in unambiguous signaling of the AD content. While conceptually simple, this method is challenging for broadcast workflows, as encoders do not dynamically change headers. Content selection is also challenging, as the value is embedded in the audio stream, requiring decoding of the asset to examine. Nonetheless, this method is of interest and being pursued.

Fast forward to modern Internet streaming, which is also standardized by MPEG through MPEG-DASH, which includes a similar audio type descriptor entitled role. Other streaming standards exist, such as the Internet Engineering Task Force (IETF) standardized HLS format, a segmented transport stream format that mirrors the descriptors identified in the MPEG system specification ISO13818-1.

### 4.2. Accessibility is About More Than Just Content Signaling

Disambiguation of accessible content audio will be a great benefit to viewers with interest and disinterest in audio descriptions. Nothing is more impactful to the interested parties than more quantity and quality of content. The tallies of content shown in Table 1 and Table 2 include children's shows, movies, and episodic series. However, live content, including news, sports and events, such as awards shows and concerts, is absent. This is partially due to the lack of regulation around live or the explicit exemption of live programs and programmers, and the technical challenges of creating descriptions in a live workflow. Describing live events is possible using a commentator calling a live baseball game. Several events have been live described, including parts of the 2018 Winter Olympics in PyeongChang, and the GlobalCitizen One World Together at Home Concert in April of 2020. These events were challenging only due to their non-standard inclusion of audio descriptions. Several professional agencies, including the Media Access Group at WGBH, have demonstrated proficiency in live descriptions. An important aspect of the creation is appropriately "ducking" the program audio during narration. A depiction of ducking is shown in Figure 1.



Lower waveform - recording audio description  
 Upper waveform is base music and effects  
 Dotted white line in upper waveform is result of ducking

**Figure 1. Audio Ducking Example**

Narrators are proficient with the timing of their descriptions, and automatic or manual equipment is used to ramp down the main program audio and ramp it up prior to and after the audio description segment. To date, live descriptions have resulted in time-delayed versions of the events, as the live descriptions were conformed to the event after completion. What will it take to produce described live events and deliver live?

As you will learn in the next section, new technologies are on the horizon to make live event descriptions possible.

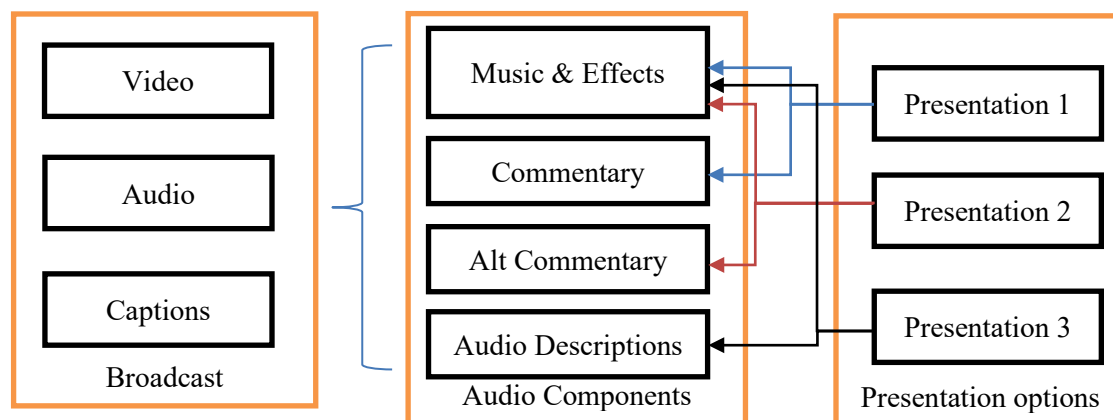
### 4.3. The Future of Over-the-Air Broadcasting

Today's over-the-air broadcast format is three decades old. Three decades ago, the Internet was accessed through dial-up modems, smartphones were ten years away, and compact discs were in their prime. The Americans with Disabilities Act was brand new at the time, and the 21<sup>st</sup> Century Communications and Video Accessibility Act was still twenty years away. So much has changed in thirty years. In 2019, the NextGen TV brand was created, which embodies the new ATSC 3.0 broadcasting standard, which was authorized by the FCC in 2017. ATSC 3.0's key features include higher resolution, dynamic range, IP delivery, interactivity and greater audio fidelity. The very modern HEVC video compression and a pair of audio compression standards, AC-4 and MPEG-H, have been selected. Selection of the national audio format is per country. The U.S. has selected Dolby's AC-4 compression as standard for NextGen TV. Two generations ahead of current over-the-air television audio compression, AC-4 is approximately three times more efficient, from a pure compression perspective, and when combined with its novel capabilities to deliver audio in substreams and objects, the bandwidth efficiencies are much higher. In the next section we will discuss some of these new audio features. [9]

### 4.4. M&E+D

While sounding like a math formula or SMS shortcut, M&E+D refers to a common Music and Effects mix and additive Dialog track. This is a key enabler for next generation accessible audio experiences. Allowing a single common Music and Effects mix to be delivered with synchronous but separate (and multiple) additive Dialog streams, M&E+D offers personalization, bandwidth savings and reduced production complexity. No longer is a complete main -- a mix of all music, effect and dialog -- required

for each audio service. The viewer/listener may select from multiple presentations, mixes of common plus one or more dialog tracks. Furthermore, these mixes can be adjusted to the listener's preference, allowing dialog to be enhanced or diminished quite easily. In the use case of audio descriptions, the narrative can be created and delivered without requiring a mix-down to the primary audio. This reduces time to deliver and offers the opportunity to lay narrative on top of a surround base. Spatializing audio is of great benefit to the vision impaired community, as it provides direction and space to the experience. Today, almost all broadcast with descriptions are delivered as monaural audio. Figure 2 illustrates how multiple audio presentations can be assembled through an AC-4 bitstream.



**Figure 2. Multiple Audio Presentations Through AC-4 Substreams**

AC-4 audio can be frame-aligned with video frames, unlike preceding audio formats. This offers seamless switching at content insertion (i.e. advertising) boundaries. Seamless switching can ease the addition of accessible advertising, ads with audio descriptions, a benefit to the viewer and provider alike.

#### 4.5. Why is this important

To understand the impact of AD, a survey of 626 visually impaired adults was conducted in 2017 by the American Foundation for the Blind (AFB). It found that over half reported watching television 4 or more hours a day -- and 65% reporting problems looking up what is on TV. Less than half were aware of assistive technologies such as AD.

Comments filed with the FCC by the AFB in support of expanding coverage of audio descriptions included the following from Jerrell Harris of Paragould, AR, which really sums up the impact of AD:

“I do not currently receive audio description from ANY of my current local channels or cable channels. I do get some audio description on NETFLIX and OMG, I ABSOLUTELY LOVE IT!!!! It changes the way I watch TV and now I feel like I am included and can SEE what’s going on, on the screen. There is SO much that we visually [interpret] on the TV screen and audio description is AMAZING and describes in detail what is on the screen and what is going on. As a blind person, it is frustrating to me to watch TV without audio description because we miss so much of what is going on. My husband will tell me what is happening on a show, but I know that gets annoying to him at times. The audio description lets us enjoy TV again, and gives a sense of confidence that we are almost on the same playing field as the sighted now because we actually have the description of what is going on the screen.”

This quotation speaks to the expansion of availability and ease of selection of audio descriptions in broadcast and cable television.

## 4.6. Automating Audio Description

Machine-based transcriptions of program audio to create closed captions is commonplace and steadily improving as models evolve and model training continues. Is it possible to automate the creation of audio descriptions? Image classification is mature and implemented in accessible applications, such as Microsoft's Seeing AI, benefitting the visually impaired as they navigate unfamiliar places and things such as food packaging. The Massachusetts Institute of Technology recently completed research analyzing television images to determine the point of focus where the actor's attention is placed [10]. Constantly describing each scene would be distracting, which makes the remaining challenge to describe the instrumental parts of a scene. It seems within reach to train domain-specific programs, such as sports or news, to machine-describe the segments and plays. Consider the ease with which a sports radio commentator describes the play of game.

## 5. Beyond Audio Description

Audio is instrumental in creating an accessible television experience beyond audio description. The FCC requires that all features of apparatus that consumers use to display television programming are accessible. Voice guidance is increasingly the means of complying with this requirement. Many features of televisions appear on video menus, which can challenge those with visual impairments. Voicing out the elements of a screen menu as they are navigated is a standard feature of screen readers applications on PCs and voice out features of smart phones. Smart televisions and cable receivers are increasingly implementing voice guidance to voice out menus and program guides. These applications typically rely on an active Internet connection to the device, which is a challenge as many smart televisions remain unconnected to networks. Embedded text-to-speech technology has evolved in performance and compactness and will increase in availability in consumer and subscriber devices.

Emergency information is new to ATSC 3.0. This information differs from Emergency Alerts, which may take over the audio and or video programming during emergency events. Emergency information may include warnings, such as weather and school closings, but doesn't automatically preempt programming. Through the use of text-to-speech technology and signaling in the AC-4 audio stream as type E, these textual alerts can be converted to audio and added as a presentation option for viewers who prefer audio narration of emergency information.

## 6. Conclusion

Cable television is accessible. That said, Internet streaming media providers offer more content and better fidelity of accessible audio than over-the-air networks and cable channels. Next generation broadcasting and pending regulation will offer expanded coverage of accessible audio on cable television programs and provide a path to more innovation in terms of immersive audio experiences, multi-language, multi-cultural accessible audio experiences and description availability on live and event-based programming. Improved signaling of audio content can improve the ability to find accessible audio. ATSC3.0 includes a new audio subsystem, AC-4, that offers presentation options that can improve personalization and lead to more descriptions on live programming.

## Abbreviations

AC-4	audio compression 4 <sup>th</sup> generation (Dolby)
AD	audio description

AFB	American Foundation for the Blind
ATSC	Advanced Television Systems Committee
CTA	Consumer Technology Association
DASH	Dynamic Adaptive Streaming over HTTP
FCC	Federal Communications Commission
FM	frequency modulation
HEVC	high efficiency video coding
HTTP	Hypertext Transfer Protocol
IETF	Internet Engineering Task Force
IP	Internet protocol
ISO	International Standards Organization
M&E	music and effects
MPEG	Moving Pictures Experts Group
MVPD	multichannel video programming distributor
NPRM	Notice of Proposed Rulemaking
SAP	secondary audio programming
VSB	vestigial sideband
WGBH	A public media producer based in Boston, MA

## Bibliography & References

- [1] ATSC Standard: A/342 Part 2, AC-4 System
- [2] Vision Health Initiative – Centers for Disease Control and Prevention - <https://www.cdc.gov/visionhealth/basics/ced/fastfacts.htm>
- [3] Disability – World Health Organization - <https://www.who.int/health-topics/disability>
- [4] C.F.R. Title 47 Chapter 1 Subchapter C Part 79 – Electronic Code of Federal Regulations – Accessibility of Video Programming
- [5] Notice of Proposed Rule Making Comments on Expanding Video Description Requirements to Increase Programming Accessibility to Blind and Visually Impaired Americans – FCC Filing MB 11-43 and comment filing 106223000511985
- [6] The Audio Description Project – American Council for the Blind - <https://acb.org/adp/appletvad.html>
- [7] ISO/IEC 13818-1 – Information technology – Generic coding of moving pictures and associated audio information: Systems
- [8] ISO/IEC 639-1 – Codes for the representation of languages
- [9] ATSC 3.0 Dolby Audio Handbook – February 2020
- [10] MIT Gaze 360 Project - <http://gaze360.csail.mit.edu>



# **A Proactive Network Management Scheme for Mid-split Deployment**

A Technical Paper prepared for SCTE•ISBE by

## **Lei Zhou**

Network Architecture Engineer  
Comcast Cable  
1800 Arch Street, Philadelphia, PA 19103  
+1 (267) 966-5010  
lei\_zhou2@cable.comcast.com

## **Robert Thompson**

Director, Network Architecture Engineer  
Comcast Cable  
1800 Arch Street, Philadelphia, PA 19103  
+1 (215) 286-7378  
Robert\_Thompson6@cable.comcast.com

## **Robert Howald**

Fellow  
Comcast Cable  
1800 Arch Street, Philadelphia, PA 19103  
+1 (215) 286-8037  
Robert\_Howald@cable.comcast.com

## **John Chrostowski**

Executive Director, NGAN Access Eng  
Comcast Cable  
1800 Arch Street, Philadelphia, PA 19103  
+1 (267) 260-3695  
John\_Chrostowski@cable.comcast.com

## **Daniel Rice**

VP, HFC Architecture  
Comcast Cable  
1800 Arch Street, Philadelphia, PA 19103  
+1 (720) 512-3730  
Daniel\_Rice4@cable.comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
1. Introduction.....	3
1.1. Benefits to Mid-split Deployment.....	3
1.2. Challenges to Mid-split Deployment.....	4
2. Problem Statement – Coexistence of Mid-split and Standard-split CPE .....	5
3. A Proposed Method .....	7
4. Case Studies .....	9
4.1. Upstream Trial.....	10
4.2. Mid-split Spurious Emission Impairments to Video STB SC-QAM.....	10
4.3. 5 <sup>th</sup> -Upstream Trial (In Progress).....	11
5. Conclusion .....	11
Abbreviations.....	12
Bibliography & References .....	12

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Coexistence MS-CM and Pre-DOCSIS CPE .....	6
Figure 2 - RFC and RHM Procedure for Mid-Split Deployment.....	9

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Results of the Eight Upstream Trial .....	10
Table 2 - Sample RHM Data .....	11

# 1. Introduction

Deployment of an expanding range of upstream split options for the return path in coax systems has many challenges. Among these challenges are old infrastructure components such as drop-amps and splitters, legacy customer premises equipment (CPE) and services sharing common or overlapping spectrum. This technical report presents a scheme that leverages remote feature control (RFC) and remote health monitor (RHM). These systems selectively enable enhanced return path high-speed internet service (HSI) for Cable Modems based on a quality of service (QoS) measurement from all CPE in the household. Specifically, current use cases will describe how DOCSIS® TLV84 is used to remotely enable mid and high-split. At the same time, downstream and upstream performance metrics, such as signal-to-noise (SNR), modulation error ratio (MER), and other metrics are remotely monitored from all CPE devices within the home and, in the case of high split and full duplex (FDX), within neighboring homes as well, to evaluate potential disruption of revenue generating services. Households are scored to determine whether they are capable of self-installation of enhanced HSI services. The advantage of this technique includes allowing: 1) scalable individualized and progressive HSI deployment with a remediation strategy focused only on customer networks where potential issues exist and 2) proactive and adaptive network operations in accordance to a varying environment thus minimizing trouble calls, truck rolls, and customer contact.

## 1.1. Benefits to Mid-split Deployment

DOCSIS has been a frequency-division duplex (FDD) access scheme, in that the upstream and downstream transmissions occupy different bands of the spectrum (Cablelabs, 2009) (Cablelabs, 2017) (Cablelabs, 2019). Standard-split operates in the 5 – 42 MHz band for the upstream transmission and has been deployed in many operator networks. Standard split has served operators well, enabling them to provide their customers with up to 35 Mbps upstream capacity since the late 1990s.

CableLabs has provided multiple options for enhancing upstream capacity with multiple versions of DOCSIS. DOCSIS 3.0 introduced the 85 MHz upstream option and with it, the ability to increase upstream capacity via channel bonding. DOCSIS 3.1 required upper edges of the upstream band to include 85 MHz and 204 MHz, so named mid-split and high-split, respectively. DOCSIS®4.0 introduces full-duplex DOCSIS (FDX) that allows upstream and downstream transmissions to share the same spectrum band where the new upstream band edge requirement extends to 684 MHz (Cablelabs, 2020).

The mid-split scheme doubles the upstream bandwidth of the standard-split scheme, which immediately translates to an augmented capacity and increased quality of service. Since many operators limit their use of the standard split band to primarily the upper two-thirds of the band, the midsplit represents nearly a 3x increase in useable bandwidth. Usually, four 6.4 MHz SC-QAM upstream channels are configured in the 5 – 42 MHz regopm, providing 122 Mbps data rate with 64-QAM modulation. The mid-split scheme is able to add four more SC-QAM channels and offers a total upstream data rate of about 250 Mbps. Alternatively, it allows configurations of advanced OFDMA upstream channels. For instance, a 48 MHz wide OFDMA channel will enable a 500 Mbps upstream data rate with 2048-QAM modulation. With another bandwidth doubling offered by the high-split scheme, more than 1 Gbps of upstream data rate can be achieved.

With all of these options available to operators, it is likely that the upstream will evolve from standard split, to mid-split, to high-split, and beyond. The good news about this approach is that all of the lessons learned from incrementally increasing capacity with customer demand can be applied to the following evolutionary step. For example, the challenges for midsplit that are discussed in this paper will also be challenges for high split, and so on. Thus, solving these challenges now with new innovative approaches

sets operators up for the future, with techniques and processes for deployments of even higher capacity upstream.

## **1.2. Challenges to Mid-split Deployment**

The relatively limited upstream bandwidth of the standard-split has been satisfactorily serving customer needs of internet access, from web surfing to video streaming. Interactivity-intense applications, including gaming, video sharing and teleconferencing have become increasingly popular, especially since the COVID-19 pandemic and work-from-home have become a new normal for many customer households, resulting in higher demand for new upstream bandwidth. This trend has accelerated changing the DOCSIS network to the mid- or high-split.

Deployment of the mid-split involves re-allocation of the spectrum used by many existing services. For example, video services on standard EIA channels 2 to 6 with carrier frequencies 57 – 87 MHz will need to be moved to make space for mid-split upstream channels. The cable plant also needs upgrades of the diplexers used in active devices, including nodes, line extenders, trunk amplifiers, and even in-home drop amplifiers. These products must support specifications for mid-split between the forward and return bands. The above requirements are challenges in and of themselves. However, the thorniest problem is in the customer premise, where a mid-split capable gateway and legacy standard-split set-top boxes (STBs) need to seamlessly coexist. The former may cause interference to the latter due to adjacent channel interference (ACI) susceptibility.

Standard-split customer premise equipment (SS-CPE) includes video set-top boxes (STBs), pre-DOCSIS 3.0 cable modems (CMs) and DOCSIS 3.0 CMs with fixed standard-split diplex filters. Mid-split CPE (MS-CPE) are usually designed with software-selectable diplex filters which can switch between the standard-split and mid-split modes. Within a customer premise, CPE devices are usually connected to the cable feed off of a splitter -- which, without sufficient port-to-port isolation, will allow the MS-CPE's upstream transmission to interfere with the SS-CPE. The "mid-band" part of the upstream signals from the MS-CPE leak through the splitter into the downstream receiver of the SS-CPE, unfiltered. Even though spectrum re-allocations are implemented, such that the SS-CPE are not expecting any services in the mid-band, thus its tuners are not tuned to any frequencies in that band, the leaked signals unfortunately raise the noise floor at the radio frequency (RF) mixers and degrade the receiving SNR for all downstream services. The consequences will be increased receiving errors and deteriorating service quality. Adjacent channel interference (ACI) susceptibility is a term commonly used for this type of indirect interference scenario.

It may appear straightforward to solve this SS-CPE and MS-CPE coexistence problem by simply replacing the splitters with ones with higher isolation specifications. This solution may turn out to be prohibitively expensive, as it may not be an easy customer self-installation procedure, i.e. truck-rolls may be needed, especially in the cases where there are many home devices, or when drop amplifiers are used. It is also unnecessary to blindly replace the splitters in all customer premises, as only a small portion of them would likely have experience-impacting effects due to the interference.

When systems allow, increasing the downstream transmission power may also help mitigate this ACI problem. However, most cable operators will have already maxed out downstream RF output power for other reasons, for instance to achieve household-per-node efficiencies. Therefore, increasing node output power is not always practical, and usually requires a whole-plant calibration for stable operations.


Deployment of the mid-split at the headend, in the cable plant and at the customer premise should be a gradual procedure. A progressive and adaptive approach is one in which the mid-split mode is “turned on” at customer premises, individually, and based on their unique site conditions, which may be more efficient and beneficial from an operator’s perspective. Modern IP-based CPEs are mostly capable of remote feature control (RFC) and remote health monitoring (RHM). Enabling protocols such as SNMP and TR-069 will allow network operators to selectively turn on certain services, based on viability, and continuously monitor the quality of services, so as to proactively adjust operation modes and class of services.

This paper presents a scheme for mid-split deployment, which focuses on the SS-CPE and MS-CPE coexistence problem. Applying a number of existing RFC and RHM technologies, this procedure includes a proactive on-line evaluation of network conditions and service quality, thus is able to isolate sites suffering the coexistence interference without requiring an installer to be on-site. The procedure subsequently supports a progressive mid-split turn-on in a home-individualized, and, most importantly, cost-effective manner.

The rest of the paper is organized as follows: Section 2 and 3 will describe in detail the SS-CPE and MS-CPE coexistence problem and the proposed procedure of online detection of it. Section 4 will present data from a field trial and other similar use case scenarios.

## **2. Problem Statement – Coexistence of Mid-split and Standard-split CPE**

To better appreciate the problem, refer to Figure 1 for an illustration of a typical configuration in a customer premise with coexisting MS-CPE and SS-CPE. The MS-CPE and SS-CPE share the cable feed through a splitter. The MS-CPE’s upstream operates in the 5-85 MHz band, while the SS-CPE’s diplexer cuts off 54 MHz and above for downstream traffic. With imperfect isolation of the splitter’s output ports, the MS-CPE’s upstream signal may leak into the SS-CPE’s downstream RF front end. Even when careful spectrum arrangement avoids the 54-85 MHz band being used by any services for the SS-CPE, the leaked signal from the MS-CPE upstream may result in increased noise floors at the SS-CPE’s demodulator, which consequently causes the SNR to deteriorate, bit error rate/modulation error ratio (BER/MER) to degrade, and ultimately service quality to be impaired.



**Figure 1 - Coexistence MS-CM and Pre-DOCSIS CPE**

A simple estimate of a dominant noise floor increment could be as follows: The maximum transmission power of a 6.4 MHz DOCSIS 3.x SC-QAM upstream is 51 dBmV per 5.12 MHz; note that the maximum modulation bandwidth is 5.12 MHz. Assume the splitter port-to-port isolation is of a typical value of 30 dB. Then, the mid-band upstream signals from the MS-CPE will generate  $51 - 30 = 21$  dBmV per 5.12 MHz interference at the SS-CPE's downstream receiver. If the downstream bandwidth is of nominal value 800 MHz, the interference will add a noise of level  $21 - 10\log_{10}\left(\frac{800}{5.12}\right) \approx -1.9$  dBmV per MHz. For a 6 MHz QAM channel, the increment of the noise power at the tuner would be  $-2 + 10\log_{10}6 \approx 5.8$  dBmV.

Assume that the SS-CPE requires a minimum of 30 dB receiving SNR for satisfactory quality of services; also assume the cable plant is calibrated such that the CPE receiving SNR is of mean 40 dB and variance  $3^2$ . With these conservative parameters, one could expect that about 10% of the SS-CPE population will suffer impairment to quality of services due to the 5.8 dBmV SNR degradation.

The interference to SS-CPE may not necessarily come from the MS-CPE's upstream transmissions directly. Even when there are no upstream channels configured within the mid-band of 42 – 85 MHz, the spurious emission of the existing upstream transmissions in the band of 5 – 42 MHz may still fall in the mid-band and then leak to the receiver of the SS-CPE, in the form of interference.

To target the root cause of the described problem, a simple solution may be to replace the splitter with one of higher port-to-port isolation specifications. For example, if a splitter of 40 dB port-to-port isolation is used, in the above calculation, the SNR degradation will become practically negligible. However, replacing splitters in customer premises may not easily be a self-installation procedure. Complex on-site work by technicians may be needed when the splitters are in hard-to-reach places. The solution will become prohibitively expensive if it has to be executed for every customer premise. It is also very inefficient based on the estimated percentage of the affected population.

### 3. A Proposed Method

An efficient solution should be proactive and individualistic. Proactivity means that the RF and network performance at the customer premises should be measured on site and in real time. Individualism means that the mid-split mode should be turned on or off for each customer premise individually, based on its unique RF and network conditions. The proposed method employs RHM and RFC technology to decide the mid-split readiness per customer premise and make a progressive deployment.

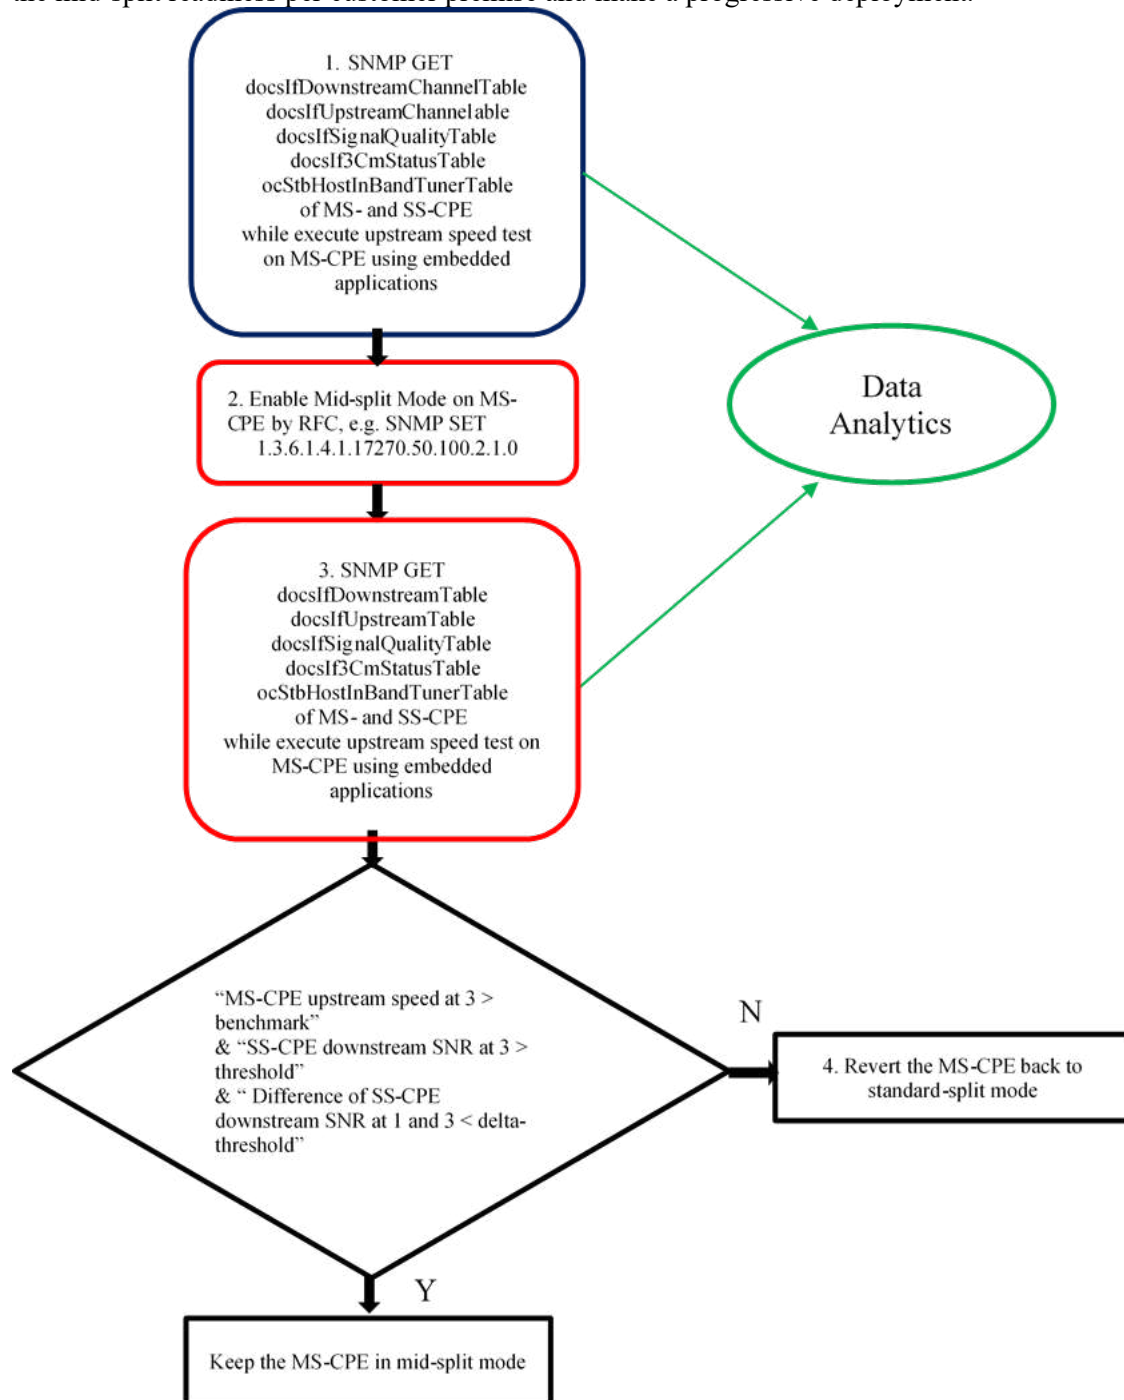


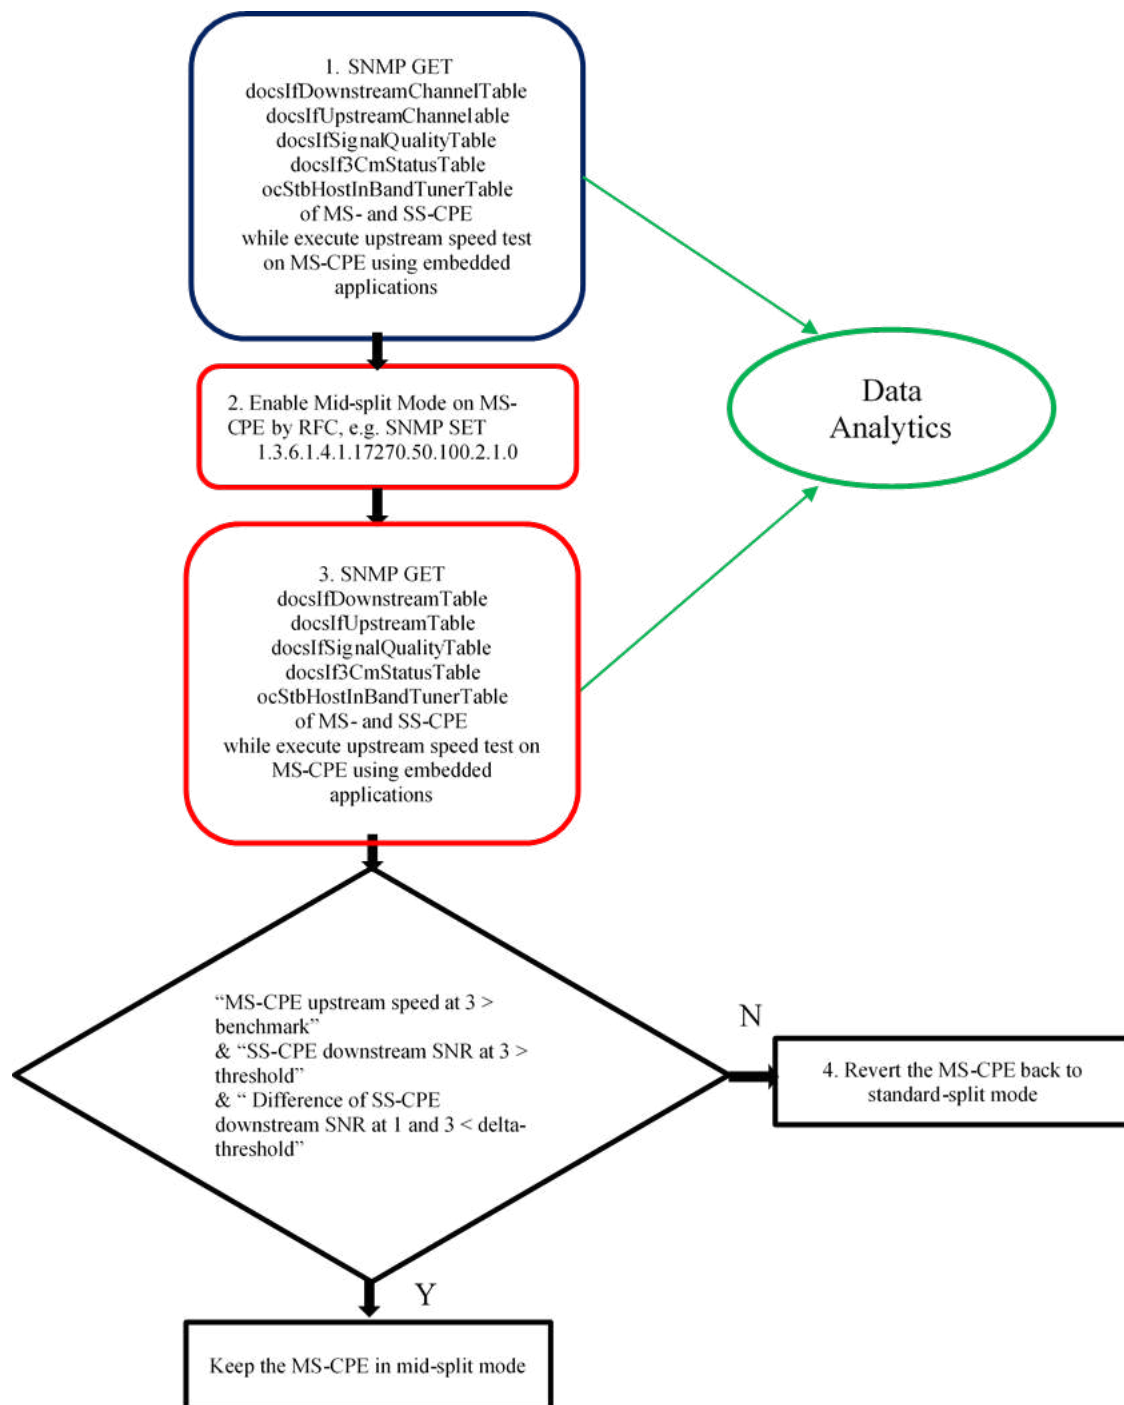
Figure 2 depicts an abstracted procedure of the proposed method. The first step is using SNMP to get the upstream and downstream metrics for the MS-CPE and SS-CPE, particularly under controlled upstream traffic that have been triggered using a speed test application. The data obtained at this step serves as a baseline for evaluating the improvements and impacts of a mid-split spectral allocation. The second step is turning on mid-split mode on the MS-CPE. This can also be achieved through SNMP, by setting the MIB object 1.3.6.1.4.1.17270.50.100.2.1.0. Then the third step is to repeat the first. Finally, by comparing the data from the third and the first steps, operators can determine the mid-split readiness of the studied customer premises. If all criteria are met, the site can stay in mid-split mode; otherwise, it could be reverted back to standard-split mode.

The speed test part of the procedure serves two purposes: It provides a direct measure of the performance merits of the bandwidth augmentation brought about by the mid-split. It also facilitates the evaluation of impairments to SS-CPE by MS-CPE as an emulated interference.

The above procedure applies to each customer premises. The data from a large number of customer premises, correlated with MAC domains and other geographic information, can offer an inference of the network readiness for mid-splits at a larger scale. Potential issues in network components, such as standard split drop amplifiers blocking midsplit upstream transmissions, could be isolated.

Note that the above procedure exemplifies a usage of some standard RHM technology, such as the DOCSIS and OpenCable MIBs (Cables, 2020) (Cablelabs, 2013). Other technologies, such as TR-069 (Broadband Forum, 2018), are equally applicable, if corresponding data models are supported by the CPE. The speed test application is generally proprietary, which, nevertheless, is widely embedded in CPE firmware. Attention must be paid on running speed tests in that they may affect customer experience; therefore, the data collection steps are better performed during scheduled maintenance windows.





**Figure 2 - RFC and RHM Procedure for Mid-Split Deployment**

## 4. Case Studies

This section presents several trial cases of the application of the proposed proactive network management method in mid-split deployment at various levels.

## 4.1. Upstream Trial

This was a full-fledge mid-split deployment trial. The RHM had focused on noise level increases in the downstream band in the coexisting scenario. The docsIfSigQTable MIBs of the SS-CPE were specifically used as a sampling of the downstream spectrum at the 8 – 24 downstream frequency points. It was predetermined that if the MS-CPE would fall into a partial-service mode (indicating standard split drop amplifier issues), or the SS-CPE downstream SNR would drop below some threshold level (indicating splitter isolation or ACI susceptibility issue), then the customer premises would be failed for the mid-band mode.

Sixty-three customer premises were selected for the trial. Automated RHM and RFC activities take about 30 minutes. The trial results are shown in Table 1.

**Table 1 - Results of the Eight Upstream Trial**

Results	Automated Testing
Pass (SIK)	65% (41 subs)
Fail (Drop Amp)	29% (18 subs)
Fail (Isolation)	6% (4 subs)

## 4.2. Mid-split Spurious Emission Impairments to Video STB SC-QAM

The objective of this trial was not to add upstream channels in the 42 – 85 MHz band, but only to enable mid-split mode on the MS-CPE in standard split plant. Its purpose was to evaluate the interference resulting from spurious emissions of the MS-CPE to the coexisting SS-CPE. Note that even the MS-CPE's transmission in the 5-42 MHz band may generate spurs in 42-85 MHz band, which leak to the SS-CPE's receiver. Therefore, the method described in the previous section is readily applicable.

The MS-CPE involved in the trial were DOCSIS 3.1 cable modems, and the coexisting SS-CPE were set-top boxes. The metrics of interest in this trial were primarily the quality impairments of QAM videos on EIA channels 2-6. The MIB ocStbHostInBandTunerSNRValue had been specifically used for RHM. To make the SS-CPE tune to the designed channel between 2-6, a proprietary remote tune application was also used.

Twenty-one geographically-dispersed customer premises were selected. Automated RHM and RFC activities took about 45 minutes. Twenty premises showed no SNR degradation, and one suffered about a 2 dB SNR drop. Sample data of passed premises and failures are presented in Table 2.

**Table 2 - Sample RHM Data**

MS-CPE																				
US frequency/Tx Power in stdsplit mode	36500000	45.5	30100000	45.8	23700000	45.8	17300000	45												
	57000000				63000000				69000000				81000000							
SS-CPE1 Tuner SNR (std-split/mid-split/mid-split with traffic)	41.9	41.8	41.8		41.9	41.9	41.8		41.9	42	41.9		41.5	41.5	41.5					
SS-CPE2 Tuner SNR (std-split/mid-split/mid-split with traffic)	43.1	43.2	43		43.1	43.2	43.1		42.9	43.2	43.2		42.5	42.5	42.5					
SS-CPE3 Tuner SNR (std-split/mid-split/mid-split with traffic)	43.2	43.2	40.4		43.2	43.2	43.2		43.2	43.2	43.1		43.2	43.2	43.2					
MS-CPE US frequency/Tx Power in stdsplit mode	30100000	43.3	17300000	44	23700000	43.3	36500000	43.5												
	57000000				63000000				69000000				79000000				8.5E+07			
SS-CPE1 Tuner SNR (std-split/mid-split/mid-split with traffic)	40.5	40.3	38.4		null	null	null		41.9	41.9	40.5		42.5	42.5	40.8		42.5	42.5	42.5	
SS-CPE2 Tuner SNR (std-split/mid-split/mid-split with traffic)	38.2	38.1	36.7		39.1	39	37.4		39.8	39.8	38.4		40.1	40.1	39.6		40.1	40.1	39.9	

This implies that future CPE deployments could continue to use a switchable diplexer, but using mid and high-split switchable CPE instead of low and mid split switchable variations. This technology re-use not only minimizes the variations in future CPE products, but also maintains high levels of return on investment in CPE to enable service growth, as customers progress from 100 Mbps to 1 Gbps upstream services.

### 4.3. 5<sup>th</sup>-Upstream Trial (In Progress)

The 5<sup>th</sup>-upstream is a 3.2 MHz wide channel with center frequency at 41.3. MHz; so, part of its spectrum is above 42 MHz. The purpose of this configuration is to increase the upstream bandwidth of the SS-CPE as well as the MS-CPE. The theory of such an operation lies in that, the transient band of the SS-CPE's diplex, combined with adaptive equalizer, may provide a sufficient 3.2 MHz wide frequency-flat channel. The RHM focus is on the partial service status, as well as spurious emissions to the 57 – 87 MHz band as discussed in the previous use cases.

The authors believe the above use case scenarios represent the tip of the iceberg when it comes to using online tools for solving a variety of challenges associated with enhancing upstream capacity. There will likely be multiple opportunities to reuse the tools discussed above to solve new, but similar, challenges associated with high split and FDX.

## 5. Conclusion

This paper illustrated a proactive and adaptive network management scheme which employs RHM and RFC to allow for progressive and individualistic deployment of a wider upstream signal path to mid-split frequency allocations, as it relates to the customer premises. The scheme is generic, such that various RHM and RFC technologies can plug-and-play. Besides automatic mid-split enablement on a per-home basis, the scheme also affords data analytics for network service quality.

The cases presented in this paper are for mid-split deployment. The proposed methods are equally applicable to high-split scenarios. As the high-split scheme is deployed in the vCMTS/R-PHY architecture and with OFDMA upstream channels, additional RHM tools may be needed and be available for the proactive network management tasks.

# Abbreviations

ACI	adjacent channel interference
AP	access point
BER/MER	bit error rate / modulation error rate
bps	bits per second
CM	cable modem
CMTS	cable modem termination system
CPE	customer premises equipment
DOCSIS	Data-Over-Cable Service Interface Specifications
EIA	Electronic Industries Association
dBmV	decibels relative to one millivolt
FDD	frequency division duplex
FDX	full duplex
FEC	forward error correction
HFC	hybrid fiber-coax
HD	high definition
HSI	high-speed internet
Hz	Hertz
IP	internet protocol
ISBE	International Society of Broadband Experts
MER	modulation error ratio
MIB	management information base
MS-CPE	mid-split customer premises equipment
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
QAM	quadrature amplitude modulation
QoS	quality of service
RF	radio frequency
RFC	remote feature control
RHM	remote health monitor
SCTE	Society of Cable Telecommunications Engineers
SC-QAM	single channel quadrature amplitude modulation
SIK	self-install kit
SNMP	simple network management protocol
SNR	signal-to-noise ratio
SS-CPE	standard-split customer premise equipment
STB	set-top box
TLV	type-length-value
TR-069	technical report, 069 specification
vCMTS	virtual cable modem termination system

# Bibliography & References

[1] Cablelabs, DOCSIS 2.0 Radio Frequency Interface Specification, 2009.

- [2] Cablabs, DOCSIS 3.0 Physical Layer Specifications, 2017.
- [3] Cablelabs, DOCSIS 3.1 Physical Layer Specifications, 2019.
- [4] Cablelabs, DOCSIS 4.0 Physical Layer Specifications, 2020.
- [5] Cables, CCAP Operations Support System Interface Specification, 2020.
- [6] Cablelabs, OpenCable Home Networking MIB Specification, 2013.
- [7] Broadband Forum, TR-069 CPE WAN Management Protocol, 2018.

# **Case Study of Social Distancing on Mentorship Programs**

## **CLEAR Program Introduction**

A Technical Paper prepared for SCTE•ISBE by

**Andrew Frederick**  
Principal Engineer  
Comcast  
4100 E Dry Creek Rd, Littleton, CO 80122  
(303) 881-6103  
andrew\_frederick@comcast.com

# Table of Contents

Title	Page Number
1. Introduction.....	3
2. CLEAR Program Introduction.....	3
2.1. Acknowledgements .....	3
2.2. About the CLEAR Program .....	3
2.3. Supported Relationships .....	4
2.4. Living in the Moment .....	4
2.5. Social Capital .....	5
3. The COVID-19 Disruption .....	6
3.1. Situation Report .....	6
3.2. Novelty of Information .....	6
3.3. Turning on a Dime.....	6
4. Silver Linings.....	7
4.1. The Technology Leap .....	7
4.2. Changing Spaces .....	8
4.3. The Work/Life Balance.....	9
4.4. Bringing It Home .....	9
5. Key Conclusions.....	10
5.1. No Substitutions for the Real Deal .....	10
5.2. The Investment of Time .....	10
5.3. Making Time to Take Time.....	12
5.4. Recalibrating Expectations.....	12
5.5. Attrition .....	12
5.6. Moving Forward.....	12
Abbreviations .....	14
Bibliography & References.....	14

# 1. Introduction

CLEAR is a mentorship program that was started in Denver by engineers, for engineers. The Mentors, Mentees, and Committee Members are all volunteers giving their time back to each new year's cohort. CLEAR stands for Comcast Leadership Engineering Achievements and Relationships.

## 2. CLEAR Program Introduction

### 2.1. Acknowledgements

Simply stated, the only way that the CLEAR Program could have been created and sustained was because of the environment, culture, and leadership. There are several instances in this paper where this is presented as pretext for creating the conditions enabling the operation of this program. The author would like to acknowledge his role and experience shaping the program over the past four years to bring this perspective to the industry at large, but not without mentioning the countless others who have founded and created the initial vision of the program in 2011. Literally hundreds of people who have given their time, energy, and dedication to pursuing the craft of mentorship and giving back to their fellow coworkers. Without their output, the CLEAR Program would have died a long time ago- the reason this program has sustained itself is by listening to feedback from Mentees, Mentors, and Committee Members about how to create a more meaningful experience and curriculum. Denice Loud has effectively been the Program Lead in Denver for the past six years, and without her dedication and execution, the sessions would not be nearly as polished and professionally presented. Her attention to detail has set the bar for expectations in how our program is run. Leslie Chapman has been the Engineering Lead in Philadelphia for the past two years and has been instrumental to launching the Philly chapter in 2019. Leslie's natural abilities have suited her ideally for making leadership accessible to her cohorts and her excellence and accomplishments in her engineering background give her the perfect experience to lead and share with each year's class. The author would like to acknowledge Denice and Leslie's contributions, not just to the CLEAR Program, but also for their contributions to this paper.

### 2.2. About the CLEAR Program

The CLEAR program pairs junior engineers with more experienced engineers on a year-long mentorship experience. The Mentee has a series of classes they will take over the course of the year that focus on professional development: Intellectual Property, DiSC, EI, Presentation Skills, Personal Branding, Leadership, and Cross Functional Relationships are the foundational topics to the program, and each year the focus for special topics changes based on executive leadership direction. The format of the program through 2019 was to meet every month for in-person sessions for a day of focused learning, and inevitably when lunch and breakfast are provided, networking. When each cohort is formed, the Mentee will propose a capstone project that will be the technical focus for the year and will serve as a vehicle to incorporate all the professional development topics together. Deliberately, CLEAR requires that all projects must not be part of the Mentee's day job- they must be cross functional in nature. This opens the door to creating a web of relationships across the organizations that don't typically happen on their own. It teaches the Mentees that the world is what they can create of it, and it empowers them to put pencil to paper, to make their vision a reality.

As the cohort moves to graduation and completing the program, the Mentee will have a demonstrable project to present and show off to their colleagues and to senior leadership at a project fair. The project fair is *THE* opportunity for the Mentees to demonstrate all the soft skills they have learned throughout the year as they showcase their technical projects as they communicate across all levels of engineering peers, management, and executives in the same conversation!



Over the past seven years, the program has been growing and one of the keys to that success has been to keep evolving the curriculum and be willing to throw things out that aren't working or are no longer relevant. The CLEAR Committee has a strong culture of adaptation and evolving with the times and is probably a key trait that has helped to carry them through the jarring transition of COVID-19. Another testament of the program's success is the launch of the Philadelphia chapter in 2019. Despite having two vastly different working environments, cultures, and locations, the core tenants of the program have been the binding agent that has kept the chapters together. CLEAR has an interesting history of how it began and how it grew to be what it is today, but that would be the subject for another discussion.

### **2.3. Supported Relationships**

CLEAR deliberately mixes oil and water together as often as possible- the Mentee project selection must be cross functional, and the partnering on projects must also be a new relationship. When the Committee pairs Mentors and Mentees together, one of the key pieces of data that's used is if either party knows each other. If there is a previous relationship there, we avoid making those matches to try and maximize the exposure of these relationships across the organization- it's a deliberate and concerted effort to force people into these relationships. One of the things that has surfaced over the years is that this works well on paper but growing these relationships doesn't just happen without some additional magic. This is where the CLEAR Committee members add that spark and give these relationships the extra push and support they need. Sometimes this pairing of "odd couples" has style conflicts, or maybe communication weaknesses that can impede this Mentor and Mentee relationship. The Committee Member forms what is known as the Triad. The Triad Committee Member (TCM) is embedded in the relationship and they are invested in making sure that the relationship between Mentee and Mentor is going well. They are also there to help provide advice, input, and guidance for any questions that come up.

One of the amazing features of the program is how much the Mentees are motivated to give back to the program and come back and serve on the Committee, or if they end up promoted to a requisite title, back as Mentors in the program! This feedback loop channels the experiences of the previous year forward into the next cohort and gives each Committee a chance to improve the program. Ideas and change are democratized as they pass through each chapter's Committee and are ultimately approved by chapter leaders in unison for the year. It's important that each chapter holds each other accountable and is supportive of moving together as a team – after all, each chapter has benefitted from the new relationships created on both ends. The hope is that each Mentee's journey in CLEAR will be unique to the individual, but the overall experience in the program will be the common bond that will be the common ground to unify Mentees, Mentors, and Committee Members from different generations of the program.

### **2.4. Living in the Moment**

Prior to the year 2020, CLEAR was an entirely in-person experience. The monthly sessions, the lunches, the project fairs- these are all designed specifically to maximize and to deliberately focus energy and resources on networking and professional relationships. The ground rules and tone that the Committee sets for the year is that the expectation is to have all electronic distractions closed and away, and to be solely focused on being in the moment to fully experience the topic and learning for the day. The expectation is that if you're in class, you have delegated your responsibilities for the day, and you are provided operational cover and support from your direct manager to be absent for the day. The goal is to create an environment and atmosphere where all the road blocks are out of the way to eliminate distractions. With this kind of environment created it was important to curate this culture and exemplify it at all levels. We created the conditions where we could maximize the time we spent together, and with all these people in the room on the same frequency it naturally brings its own energy. Running a session was no picnic, but the energy quickly overflows once we get into breaks and have lunches together. It

was a symphony of managing and releasing energy in a deliberate setting. These sessions were as much about the social growth of the Mentees and Mentors and Committee- the class and learning for the day is a convenient means to package some truly important content into a setting which promotes relationship building. A few years ago the logo was redesigned, and the Committees chose a design that emphasized the “R” being represented in a different color. The reason behind this choice is because the Relationships are the most important product from the program. It became apparent over the years that the projects, ideas, and technology ebb and flow, but the people seem to be the biggest constant in a healthy organization. The relationships are what seemed to outlast any idea in a digital world and are what brought many Mentees future opportunities after graduating from the program.

## 2.5. Social Capital

The emerging theme here is that for this kind of mentorship program to be successful, it had to invest heavily in social capital. Long before this pandemic, Comcast has been investing in social capital and empowering their employees to do the same. Mentees who graduate the CLEAR Program return to their teams and get a chance to flex their new skills and experiment with what they have been learning. Alumni know how to network and aren’t afraid to reach across organizational boundaries, and their experience empowers them to communicate with leadership when they have an important idea. It can take a few years to build this kind of cohesion and camaraderie between the Mentors and Committee members, but once this trust is built, the group can go far and fast together. Maggie Heffernan says it best in her TED talk,

*“When the going gets tough, and it always gets tough if you are doing breakthrough work that really matters. What people need is social support, and they need to know who to ask for help. Companies don’t have ideas; only people do. And what motivates people are the bonds and loyalty and trust they develop between each other. What matters is the mortar, not just the bricks.”*

Mentorship programs like CLEAR are essential for companies to have as part of their DNA- these programs are what gives companies the ability to weather the storms like COVID and come out the other side moving in a positive direction. CLEAR focuses on the strength of the organization, not solely on the individual. Back to Maggie:

*“And when I talked to producers of hit albums, they said, “Oh sure, we have lots of superstars in music. It’s just they don’t last very long. It’s the outstanding collaborators who enjoy the long careers, because bringing out the best in others is how they found the best in themselves.”*

CLEAR has grown to the point where their social capital compounds, even as it’s being spent. Every year there are new individuals coming back to the program to help as Mentors or as Committee Members. The main reason these groups can survive is because they find other like-minded people who enjoy giving their time to others.

### **3. The COVID-19 Disruption**

#### **3.1. Situation Report**

The events that unfolded going into 2020 may not happen again within our lifetimes, but it is important to reflect on the state we were in and how it was disturbed by a global pandemic.

The awareness of the pandemic was not beginning to be realized in the USA until the end of the first quarter of the year, and it was largely something that caught businesses off guard. Companies had work from home policies and plenty of employees contributing in this manner, but what the world wasn't prepared for was the tectonic shift to moving most of their employees to the safety of their homes. It's important to note that this transformation had to be created at rapid speeds across massive organizations, which is about the equivalent of power sliding a battleship into a parking space. The CLEAR program was on the receiving end of the information pipeline which meant the Committee had precious little time to react and pivot.

#### **3.2. Novelty of Information**

In the early stages of the pandemic, there wasn't a unified sense of purpose or mission about how our lives were going to change. It's important to note that there were voices across the spectrum which had varying degrees of opinions about what the next few months held, and how temporary or permanent these changes would be. CLEAR is a volunteer organization and it was important for the group to reach consensus quickly about coming to terms with the situation and how to best move forward.

The truth of the matter is, nobody knew exactly what was going to happen in the short or long term, so we're dealing with the novelty of information mixed in with human emotions. The plans that were made "yesterday" were thrown out the following day. Responses ranged from denial to panic depending on personal situations and how close to home the virus was hitting. Executive leadership was key in these early stages to setting the tone and direction for how we were going to respond to these changes to our business. Within the company, the response from leadership has been amazing. Experts had been consulted and leadership began providing more frequent updates, coupled with meaningful plans about how our business is going to change and walking us through this transition. This is important to mention because it was absolutely critical for our volunteer organization to have an example to follow and emulate. Leadership being comprehensive, direct, and focused allowed the message to penetrate to the core of the organization and offered a path forward. This new path was not a like-for-like replacement. It was a new landscape to navigate, fraught with a unique set of opportunities and challenges. Had the guidance from above been delayed or not effectively communicated with timely updates, the CLEAR Program would have had a difficult time trying to execute their program in 2020.

#### **3.3. Turning on a Dime**

The 2020 CLEAR program began its year as it has for the past several years by collecting our feedback from 2019, digesting it as a committee, and then having committee leads parlay that into a conversation with the other chapter and represent how we want to move the program forward for the year. Nobody in either chapter could have identified the potential threat resulting from the virus and planned the necessary changes. We got so far as to have one of the first kickoffs (the Mentor kickoff) in person in Denver. The plan for Denver was to push forward and try to host our Mentee kickoff on schedule the following week. COVID's disruption had the Philadelphia's chapter decide to inject a delay in the start of the program to allow them time to recover their footing and figure out how they were going to proceed. Bear in mind that the East Coast of the USA was starting off in a vastly different situation than the Midwest was which

influenced the different paths forward chosen by each CLEAR chapter. It turned into a favorable formula for the setup of the 2020 year, as this placed the Denver chapter in the hot seat to be the experimental group who got to trial the virtualization changes to the program first and passed along the learnings to Philadelphia so they could minimize the novel challenges we were uncovering along the way and host a better experience. This offered yet another iteration to which the formula can be evaluated and tweaked without having to wait another year to try something different. It allowed us to iterate quickly and early, and when we found something that worked, we brought those learnings forward for each chapter's benefit.

When the pandemic was hitting home, both chapters put forward discussions about what it meant to continue forward. Honestly, we entertained the idea of not hosting the program in 2020. We also entertained delaying the start of the program for a few months. At the end of the day, both chapters decided that it was best to figure out a path forward and seize the opportunity. The mindset was that we needed to choose a conservative route forward, and the most conservative path showed that this virus would possibly be impacting our lives in 2021 and beyond. If 2021 was going to be a repeat of 2020, we needed to take our lumps sooner rather than later and be leaders in this space. The Committee should be applauded for their hard work keeping the program together and staying the course in the face of such adversity!

## **4. Silver Linings**

### **4.1. The Technology Leap**

Being a technology company and having a strong culture of innovation was essential to transforming the workplace overnight. This created the spirit of challenge and curiosity to explore the question of this mentorship program's ability to survive this conversion. We explored this question in good faith and inevitably the tool that brings us all together isn't a great tool for making us feel connected. The same online environment we create for learning and sharing information in a top-down approach inevitably leaves a gap for people to socialize and interact. Bringing large groups of people together is a powerful top-down way to distribute information- this inevitably creates a fast ingestion of information but leaves no time for the social aspects of working on a team. Collectively as a society, this online format should be considered a new venue, and will influence different personality types in new ways. We're naturally not going to get a lot of lateral networking opportunities in this top-down format. To with; we must find ways to synthesize effective team building in a virtual environment. Each of the CLEAR Committees has found different ways to cope that are meaningful to them. Philadelphia, for example, gives a five-minute break before starting their calls. Denver dedicates this time for small talk and surface level conversations about how we are doing in our personal lives. It's important to start a practice now of taking some time to have this relaxed conversation in order to build a strong team that will be able to respond to new challenges.

It's tough to say, ultimately, how the cohorts for the class of 2020 will feel at the end of the year and how satisfied they will be having undertaken this experience of Mentorship in a new format. By the time of publication, the ending of the story will not be known and could be covered in a future update to this paper if there is enough interest. The concept previously discussed about creating time for networking didn't occur to the Committees until feedback from the Mentees started to come in through surveys and our monthly check-ins. In hindsight, it feels like an obvious miss that we didn't seize on this opportunity sooner in the year. For next year, the Committees are considering changing the curriculum to build the socialization events into the third month in the program after the teams are formed and capstone proposals are due. The conclusion CLEAR arrived at is that it's important to lean in early and spend time strengthening the team and interpersonal dynamics early and often. Not taking time up front when

forming the new cohort probably held Mentees back from feeling connected to the rest of the Mentors and Committee Members who have already got previous experience with the program. Another piece of feedback CLEAR has found to work for these socialization opportunities is to create small groups for these interactions. The online format doesn't necessarily scale well beyond more than a few people in a conversation at once- the upper limit is probably around five or six individuals. After a certain point, large groups are going to be counterproductive to the networking experience.

## **4.2. Changing Spaces**

Technology has opened new doors for us, but it would be worth mentioning that working from home shouldn't mean that we always feel obligated to be working from the same desk. There are some obvious limits to what's acceptable on camera – that said, it's possible to both work and be in an enjoyable space at the same time. After all, why shouldn't we make lemonade when life has handed us lemons? This doesn't work for everyone – someone must be facilitating and running the agenda. The result should be a relaxed group setting where the team can take a call in a fresh space and recharge themselves mentally. Take turns leading on different occasions so everyone can get the benefits from this practice.

Having a home office should be a consideration for every professional person or family when they acquire a new residence, post COVID. Working from the kitchen or the dining room may be effective for a few weeks or months, but it's not sustainable in the long term. These statements are made with consideration and understanding that it's not that easy for everyone to just do. But if professionals are finding themselves with a move coming up, having dedicated office space will be the equivalent of having a guest room in your home. It's probably something new owners are going to be actively looking for going forward, and reflecting on the pandemic, it's probably a good thing to have in general. COVID isn't the first or the last disease to transform the world – in the spirit of using this tragedy as a learning experience and a preparation exercise, it would benefit us all to have dedicated workspaces at home that we can fall back to should there again be conditions to keep us from socially collecting in an office together in the future.

It may be helpful for leaders to set some expectations for what a productive home office should be comprised of as employees are looking to make changes to their homes. Some things that should be standard in a home office for a technology professional would be a dedicated desk or workspace, at least one external monitor, and a door that can be closed. Ensure that the natural and artificial lighting can be controlled. Employees should have their own dedicated Internet connection with reliable Wi-Fi or direct hard line connections to their cable modem – not borrowing Wi-Fi from a neighbor. Should you have a family working and learning from home, consider upgrading your service if that's an option to support the needs of the family unit.

Most families and individuals renting or buying domiciles in 2019 didn't consider additional rooms in the off-chance that they'd be working from this space long-term. The reason is simple – homes and real estate are expensive, and it would have been cost prohibitive to try and plan for that pre-COVID. Professionals who are changing spaces in the COVID market should actively seek out situations where they can create a workspace that will suit their needs for the next few years – possibly even for the long term. Employers are also going to have to be flexible and compromise, here, too. If employees are going to be moving a little farther away from their jobs to be able to reasonably afford this extra space without assistance, we need to remain flexible with working from home policies as the situation (hopefully) changes in the future for a more traditional in-person experience. It's not convenient to buy/rent/sell homes and move families on a whim, so we should be mindful about permanently allowing work from home flexibility in the future for everyone as the commute times may become a new factor as employees move to areas that fit their new needs.

### **4.3. The Work/Life Balance**

One skill that can be sharpened like a sword during COVID is patience. Things are probably not going to go back to normal for quite some time – possibly years – possibly never. We all need to have a survivor’s mentality, and it’s important to keep this in mind to maintain a healthy work and personal life balance. This new normal has finally removed the burden of valuable time spent idling in traffic and instead gives us the opportunity to become stronger families. It’s exciting to think that this is a chance to break out of the mindset of forcing ourselves into the office every day. Just as a general note, not even specific to the CLEAR program, everyone should be practicing strong mental health routines and habits now and be prepared to settle in to this new normal. We have seen varying degrees of success with our CLEAR Mentees retaining this balance. Unfortunately, our impressions would have us say that most of the Mentees struggle with breaking away from their desks and getting outside for fresh air and exercise. This is a dangerous trend that has long-term health implications. It takes energy and active polling to reach into this level of concern for our coworkers as human beings. Those of us who are taking care of ourselves are doing their best to help break the barrier down and remind our colleagues that if we are to emerge from this situation, it’s going to take a lot more effort than usual to stay healthy and strong when so much has been taken from us.

Non-traditional meetings can be a fun way to change up the scenery and find a way to interact in a small group. All government advisories should be strictly adhered to, but within that framework there may be a comfort zone that people can live with to maybe go for a walk in a park, a hike, or on a bike ride. This is a great opportunity to refresh yourself outside and build a relationship with someone in a safe manner. This may be more difficult in urban settings but look for creative ways to use your lunch hour or set a 7:00 a.m. meeting and go for a hike before work. COVID has created a mentality that it’s OK or expected that you’re working from home constantly. Most employees surveyed during this pandemic feel like they are trying to outcompete their peers and that they are literally competing to keep their jobs. Especially for newcomers to this landscape, it’s important to hear that within reason, you are empowered to spend time in other meaningful ways.

It’s important to remember not to burn out. We all must find ways to stay sane, and in case it wasn’t clear before, COVID has given us all a fresh start on creating new routines and changing how we spend time in our lives. Start now, not in year three of the pandemic. Protect your physical and mental health – without these assets, it will be impossible for us to sustain our families, teams, and companies through this change. Use this as a chance to examine how you are managing your energy. Ask questions of yourself about if you’re being stretched too thin, or if you don’t have enough on your plate. Be honest – COVID has changed our home family dynamic and upended our lives. If your home life is drawing more of your time and attention, ask yourself what you do in your professional life that you can give up and get some of the time back. There will always be new ways for you to contribute and step up into new roles – we must remember to take care of ourselves first before we can share our best with others.

### **4.4. Bringing It Home**

Some of the blessings of being forced to work from home is being able to see leadership in their home setting. Comcast leadership has opened themselves up and brought their employees into their homes and provided us with meaningful and timely conversations about how we’re moving forward as a company. It’s been a connective experience to be in such an intimate setting and it shows great leadership to be vulnerable. It’s effective within the Comcast culture, and it’s working. In a post-COVID world, we should leave the door open for these settings to continue to be experienced by our employees. In the meantime, it continues to motivate and set the tone for how we move forward. It’s humanizing to see that your favorite executive has a dog that can also interrupt meetings when the Amazon delivery driver shows

up, or that another leader is into restoring an old VW bus. These are truly memorable distractions that help bring us closer together and feel like a family.

## **5. Key Conclusions**

### **5.1. No Substitutions for the Real Deal**

One of the hard realities that must be acknowledged is that trying to create this kind of mentorship program during a pandemic may be impossible or may be an exercise of diminishing returns. CLEAR had the benefit of growing in popularity over the past seven years and was able to build this culture of creating meaningful relationships and embed it into the Committee and Mentors. The CLEAR Program has concluded that there are no real substitutions for building that social capital in the electronic space, or perhaps the working culture is not yet there. The broad population is burning out from being on camera all the time. Over time this disruption may become more accepted and ingrained into our experience, but many people are struggling with the new challenges present in their everyday lives, plus this new work environment. When this pandemic does finally end, it's an important reminder for leadership at all levels to remember that we need to actively create time to value each other and have meaningful social experiences that bring us together. Programs like CLEAR will likely sustain themselves over the next few years because they have a strong culture built around their nucleus and can probably weather the storm. When society finds a way to safely come back to the workplace together, we need to embrace active team building and make this a priority for business. In the meantime, this will be the next challenge at every leader's feet. For existing teams, this problem may be diminished because of time spent together. But the next pandemic may be more deadly, and it's important that we embrace this opportunity to prepare our companies to survive greater threats to our businesses. Even if within our lifetimes we never have to be forced to work from home for a long duration again, we'll still have a stronger workforce that doesn't feel as constrained by virtual walls – it also opens the door to more diverse future employees who no longer have to be embedded in metropolitan centers. These teams are going to be central to creating and maintaining the technologies that run our businesses, and we owe it to them to invest heavily in their success. Virtualization challenges are not going to go away, and the longer we operate in these conditions, the more exposed we become to the challenges of starting new projects and forming new teams. The author feels like there's a lot more work to do in this space, and we may have to change our perceptions about how we're evaluated at doing our jobs and what performance measures are weighted towards.

### **5.2. The Investment of Time**

We have probably all felt this experience in 2020 that time is simultaneously going by faster and slower. The Groundhog Day cycle of getting up and being on camera day-in and day-out has created an emotional and mental crash where some have adapted, and others continue to struggle. This means that what we did before, now costs us more. There has been a big shift in the investment of our time, and ultimately, our energy. It takes significantly more energy to focus on listening and watching at the same time, plus incorporating our natural social anxieties about being on camera means that it takes a lot more energy and time to do everything with video.

Consider going 20<sup>th</sup> Century – especially if it's a meaningful relationship. A phone call can feel like something we used to do a lifetime ago. It is a bit more impersonal, but focused listening and engagement on a meaningful level in a one-on-one situation can create the same meaningful connections. The great thing about phones is that we can walk around and be free – watch a nice sunset, watch a child

play – passive viewing experiences allow us to have a little room for our minds to naturally wander while still allowing fully focused out-of-band engagement.

With this oversaturation of communications taking place in both the visual and audio formats, we should be honest and consider that not even a year ago, most of our meetings that took place remotely were done over the phone. Taking time to call someone instead of setting up a video conference can be a nice break for both parties. It's much easier to move yourself outside to get some fresh air, or take a small walk, rather than being tied to a desk or reliable Wi-Fi to hold that face-to-face conversation.

Being on camera all day at work is exhausting and creates an even more energy draining experience than we've been used to for our working lives. We should consider trying to remove some of this extra energy and go camera off with intention. Can you move non-critical meetings to audio only? Is it possible to save our camera energy for the important meetings? Leaders should be encouraged to set some intentional boundaries, so we avoid burning each other out. Maybe a "No Camera Friday" or something easy to show that it's OK to turn our cameras off, and not have to be a TV anchor for the day. It's impossible given the current working climate to assume that everyone is always going to have a perfect hair day while dealing with family and children and working from home for the 200<sup>th</sup> consecutive day in a row. Even the most consummate professionals will find that the energy must come from somewhere, and productivity is probably the first to fall to the wayside.

Another mindset we should all look at making a shift towards in this pandemic is to trim our calls down to 30 minutes at a time, 60 max for critical content. Time is a commodity in a pandemic, and there are only so many hours in a workday. It may become necessary for corporate cultures to shift away from so many meetings and allow employees half or more of their work week to dedicate towards productivity. In the always-on environment in the pandemic where we're all trying to save our jobs, it's important to reserve our energy for things we can actually produce and create. It helps us feel fulfilled in our roles and creating output for most people is a fulfilling experience. At the end of the day, it's all about carefully balancing collaboration with productivity. Perhaps the impact of COVID has shifted us into a modality where we're trying to maximize or overcorrect for collaboration in favor of productivity. It may be helpful to trim out unneeded calls from our calendars and pull back on commitments so we can self-manage this balancing act on our own accords.

CLEAR struggled with this a lot this year. Early on we took a poll to see if our cohort wanted to keep our sessions into one single day, or if we split them up over two days in the month. Denver and Philly both voted to keep their schedules into a single day, but for 2021 this decision will be revisited. CLEAR has such a dense agenda during a non-COVID year that we often found ourselves underbudgeted for time and our sessions ran long. Even when we had our sessions fully prepared and time boxed to our expectations, we found it difficult to keep the group together electronically over the course of the day. A continued theme was that everything seemed to take a bit longer when we were asking groups to collaborate. When the CLEAR Program sets their schedule for 2021, we're going to be pulling back on the content to try and consciously open up any group work to deliberately have more socialization time included. Next year's practice should embody more robust group time to allow not just focus on the learning topic, but also for the social interactions within the group.

Another learning point was not scheduling more time up-front in the year. We learned as time went on from our monthly check-ins and surveys that the networking wasn't happening well. For next year, we're going to build more time in the program for this socialization early on after the teams have been formed.



### **5.3. Making Time to Take Time**

It bears repeating again, as this is a central thesis of this paper, that time is a commodity in this pandemic. With the experiences we've seen running the CLEAR program, we've noticed that it takes a lot more energy to do the same things we used to do pre-COVID. It's hard to imagine that technology would be the barrier, but in a way, we're mismanaging ourselves into this modality. It's easy to see why, because the desire to socialize and have a more interpersonal experience will always be there. The old expression says if you're going to do something, you might as well do it right. This mindset extends to mentoring and networking – if you aren't making time to get to know someone better, what are your expectations for the relationship to grow and develop? If time is considered valuable and sacred, it places more meaning on those moments that are spent together.

### **5.4. Recalibrating Expectations**

Once committed down the path of hosting a 2020 CLEAR cohort, it was important to begin managing our expectations for the year. As a Committee, we discussed our threshold for success this year wasn't the typical 100% graduation rate we normally target, but if we got 66% of our Mentees across the finish line, we'd have done a decent job of holding the program together. As time passes and without knowing the future, both chapters are likely going to approach a 100% graduation rate during the pandemic, which is a testament to the leadership of both chapters, and the networks of professionals that are being created and sustained by this program. As chapter leaders, we felt it was important to not let our Committee get singularly focused on repeating the past seven year's performances. It was more important to move forward and make mistakes than to be distracted by metrics we would normally hold ourselves accountable to during a non-pandemic year. It was important to give ourselves that break, mentally, and lower our expectations and the measurement of success. That action relieved the pressure, allowing the committee to focus on the larger issues of restructuring the program in the virtual setting and solving for those problems. We didn't expect to get it all right this year, but it was important for us to get in the batter's box and start taking swings. There will still be several months of the program that transpire after this paper is presented – there are no guarantees that the current trajectory will hold. But we are hopeful that as an entity, we've learned enough from this year to reimagine the 2021 structure, and simultaneously having the experience for running the program virtually.

### **5.5. Attrition**

It's important to note that CLEAR wasn't immune from attrition due to the pandemic. There was a pair of Mentees who withdrew from the program in 2020 that was a direct result of the pressures placed on their personal and professional lives. As a Committee, we mentally prepared ourselves of the possibility of sustaining major losses as we transitioned to the virtual environment.

### **5.6. Moving Forward**

When the pandemic was breaking in the USA, the CLEAR program had to make a decision. The syllabus takes place over the course of the calendar year, with each cohort graduating after Thanksgiving in early December. What were we to do? 2019 was the most successful year of the program – with an eye for founding two additional chapters in 2020. We had varying perspectives about the duration and intensity of this change in work styles, and nobody had the crystal ball to give us a prediction where we'd be in four weeks, let alone four months. The Philly and Denver Committees had to discuss this and quickly reach an answer about how to move forward. Should we pause for a few months and try to start late? Or do we cancel 2020 altogether and hold out that 2021 will be back to normal? The Committee had a lot of doubts about the effectiveness of taking a 100% physical program and transforming it overnight into a

virtual one, especially without any prior thought or preparation. Most of the foundational concepts and principals for the constructs of the program centered around being there in person, distraction free, and deliberately creating time for these horizontal networks to grow. The projects are a great vehicle for the Mentees to drive on this journey, but as was discussed previously, the strong relationships were truly the key outcomes from the program. It has taken and will continue to take a lot of time, good ideas, and bad ideas to figure out what works to make up for the ground that COVID has taken from us.

The author would consider the decision to move forward under complete uncertainty to have been particularly wise choice in hindsight. One prominent theme that has emerged during COVID is that scientific and medical information keeps evolving and our company is also changing the dynamics and conditions on which we could engage each other in person. Moving CLEAR to a virtual space allowed us to move forward and face the challenges of this space head on. The author is the Engineering Lead for the Denver chapter and on the inter-chapter panel, allowing visibility to how two geographic locations within the company tackled challenges in similar spaces, how effective those problem-solving strategies played out, and the unique challenges that were surfaced from each location. Being able to have a larger sample size effectively doubled the interactions and outcomes and provided a meaningful backdrop of experiences to base the advice, guidance, and observations that are foundational to this paper. The goal of this paper was to provide a deeper view inside the CLEAR Mentorship Program, and to serve as a mirror and a reminder that even a prestigious institution like CLEAR is not immune to the challenges presenting themselves during a pandemic.

This paper has discussed the sustainment of an existing mentorship program with a historic context and rich culture. What this paper does not address is the fact that life must continue forward outside of CLEAR. Leadership will be forced to contend with company reorgs and creating new products as we have done in the past, but with new challenges and barriers. How will leaders look to solve the issues of creating new teams and build those horizontal relationships as we found new teams, or create new organizations? On paper, the exercise is simple, but the author believes that this will be a large question that's on CLEAR's horizon for years to come, as well as the leaders within a given business. The crux of the problem isn't how to move forward when we've built an amazing ship, we're navigating to a destination, and have the winds at our backs. What do we do about the boats that haven't been built, and the teams who haven't yet set to sea? CLEAR's issue of adaptation may be an opportunity and a challenge at the same time – on one hand, CLEAR had to change directions, not create something from new. It's the author's hope that many of the challenges presented with virtualization in 2020 would be points of interest for anyone starting a new team or for moving an organization forward and is hopeful that the lessons learned here will be a useful analog for how we each have to move forward.

In conclusion, the author will go out on a limb and put a prediction into print about how long it will take most of us to adjust to these conditions – barring a vaccine or other breakthrough. 2020 will feel like the year that didn't exist for most people – we're in the denial stage, still, especially in the USA. 2021 will dawn, and there will be a realization at some point where people will observe that we've been under these conditions for a year, and the reality will start to sink in – acceptance. By the time 2022 arrives, the Groundhog Day novelty has faded, and we'll be in the new normal. 2021 will see some oscillations as we struggle to cling to the past and the reality of coping and dealing with the new normal – in this year do we finally start to achieve group awareness that this problem is not temperamental and employee cultures will become more accepting to this change? We probably won't have the full focus of problem solving in this space achieved until sometime in 2021, with 2022 becoming the foundation for the new normal.

Stay healthy and safe.

## Abbreviations

CLEAR	Comcast Leadership Engineering Achievements and Relationships
EI	Emotional Intelligence
TCM	Triad Committee Member

## Bibliography & References

Leslie Chapman – CLEAR Engineering Lead, Philadelphia

Denice Loud – CLEAR Program Lead, Denver

David Eng – CLEAR Mentor

<https://www.fastcompany.com/90537949/2-science-backed-benefits-of-making-small-talk-with-coworkers>

<https://www.bbc.com/worklife/article/20200421-why-zoom-video-chats-are-so-exhausting>

[https://www.ted.com/talks/margaret\\_heffernan\\_forget\\_the\\_pecking\\_order\\_at\\_work?language=en](https://www.ted.com/talks/margaret_heffernan_forget_the_pecking_order_at_work?language=en)

# **Expediting New Product Deployments with Agile Operations and DevOps**

A Technical Paper prepared for SCTE•ISBE by

**Andrew Frederick**  
Principal Engineer  
Comcast  
4100 E Dry Creek Rd, Littleton, CO 80122  
(303) 881-6103  
andrew\_frederick@comcast.com

# Table of Contents

Title	Page Number
1. Introduction.....	3
1.1. Foreword .....	<b>Error! Bookmark not defined.</b>
1.2. Traditional CI/CD Model Overview .....	3
1.3. Multitasking and Task Changing .....	4
1.4. The Eighty Percent Rule .....	4
1.5. Hitting the Moving Target .....	5
2. Launching A New Product.....	6
2.1. Universal Truths .....	6
2.2. Challenge of Scale .....	6
2.2.1. Inventory and Accountability .....	6
2.2.2. Monitoring .....	6
2.3. Challenge of Change .....	7
2.3.1. Data In, Data Out .....	7
2.3.2. The Hardware Lifecycle .....	8
2.4. Meltdown Intel Security Vulnerability – Retrospective .....	8
2.5. Who Can Move My Cheese? .....	9
3. System Users .....	9
3.1. Someone Else’s Problem.....	9
3.2. Accessing, Viewing, and Changing Data .....	10
3.3. Why Did We Wait This Long? .....	11
3.4. Build a User Interface That is Meaningful; Build Your UI in a Meaningful Way .....	11
4. Conclusion.....	11
4.1. Key Conclusions .....	<b>Error! Bookmark not defined.</b>
Abbreviations .....	12
Bibliography & References.....	13

# 1. Introduction

Deployment Engineering in the cable sector is where the rubber meets the road. These groups are responsible for taking products from the prototype stage to a scaled solution across the entire enterprise. Deployment Engineering is an in-house specialty that creates the cloud resource experience and delivers that to the local headend. An idea begins much like how a single raindrop falls thousands of miles upstream, then gathers and runs to the ocean. This is a useful analogy for the lifecycle of a product, too. By the time it reaches the ocean – in this analogy, our customers -- it reshapes the environments that it passes through. Cable products can be based on web applications, but when dealing with real time video or streaming video, it becomes impractical to distribute content centrally, and must be decentralized in order to maximize efficiency. One can simply look to the Content Distribution Networks (CDNs) that the industry's Multiple Systems Operators (MSOs) have spent billions of dollars building to move video to the edge as evidence of this truth.

Deployment Engineering holds the unique perspective of being at the nexus of product creation and customer integration. Over the years, the methodologies change (waterfall, agile, lean, scrum, Kanban, CI/CD, etc.) to embrace new ways of tackling problem spaces, and offer meaningful ways of reinventing the wheel. Such methodologies will continue to change and evolve with time – however, the problem space of deploying products at an enterprise scale has a core foundation of challenges that are common to nearly all products. Therefore, it becomes more critical for engineering leads, architects, and executives alike to understand the environment and problem spaces than it is to be masterful of any given tool used on the same old problem. Whenever shifts occur to adopt the latest methodology, as an industry, we put our momentum at risk and can be prone to recreating and reliving the same mistakes. We sacrifice the valuable experiences and lessons learned to move to a new methodology, and in doing so tend to be more focused with how we're operating in the new system rather than staying focused on the big picture. This paper aims to reinforce the common elements within new products and deemphasize the tools used to build them. In other words, technique matters more than the actual tools used to build something.

## 1.1. Traditional CI/CD Model Overview

It's helpful to take a moment to acknowledge the landscape of contemporary software and hardware implementations in an MSO. An appreciable trend over the past decade is to build your own system, using off-the-shelf vendor hardware. In some cases, that extends to the development of in-house and at-scale technology solutions that can be syndicated to business partners. It's a high risk, high reward endeavor, and being the first to market can be the difference on that investment paying out or not. To that end, the current method of maximizing efficiency and speed is the Continuous Integration/Continuous Development (CI/CD) pipeline. Some define this acronym as Continuous Integration and Continuous Deployment. Others, to be tongue in cheek, call it Continuous Integration and Continuous Disruption. In any case, and if not managed properly, CI/CD can become a dangerous feedback loop that pulls on products like a black hole -- and can grind a product's market introduction to a halt.

Why are cloud services successful to the point of creating a \$150B per year industry? Because the cloud offers the most direct and rapid route to putting bits on the pipe. The ability to instantly scale one's computing needs is insanely attractive, but it doesn't come for free. For web services, where traffic is bursty and light, and done over a Transmission Control Protocol (TCP) connection with retries, it's relatively affordable to purchase CPU, memory, and disk in a datacenter, spin up an application, and you're off to the races. In the cable video portfolio, however, it's not always possible to create a good user experience by running applications out of a hosted datacenter, and the need to move large volumes of data can be cost prohibitive in leased space. Often the locations of national and regional data centers can be hundreds of kilometers from customers, which makes them impractical. This begets the need to build

custom hardware solutions at the edge. These products are typically homed at the edge, or use tiered caching models to distribute efficiencies of storage and transport. One of the main trappings of CI/CD is that it can be “loved to death,” meaning it can be so appealing that it becomes overused and negatively impacts the outcome. This vulnerability that presents itself with many modern methodologies, including scrum, agile, lean, etc. CI/CD must be mindful about the state of the platform and the ability to make architectural changes at scale. Without this perspective, if the sole focus is on constant delivery, it’s easy to create more work than resources can support, and accumulate unmanageable volumes of technical debt. Technical debt is usually accumulated when decisions are made that favor short-term and easy solutions, rather than using a better approach that would take longer. Tech debt, like financial debt, carries interest and makes it harder to implement changes to a product. For best results, a careful balance between change and progress must be struck that emphasizes the maximal efficiency of this model, rather than the act of change and flexibility itself.

## **1.2. Multitasking and Task Changing**

It’s no secret that there’s benefit to being able to simultaneously handle multiple workflows. Modern CPUs are a fantastic example of being able to handle large workloads and change tasks, for virtually an unlimited amount of time. A CPU doesn’t care what kind of processing it’s doing -- it marches along endlessly forever. It feels obvious to say, but human beings are not CPUs. A certain amount of task switching and busy work is helpful, but taken to extremes, it can quickly overwhelm teams of humans who aren’t able to work around the clock. To use an example, if you were to write a sentence ten times on a piece of paper (“I am a very good hexadecimal mathematician.”) using two different methods and timed this exercise. It takes a lot longer to write one letter on each of the ten lines serially, rather than writing the complete sentence ten different times. Task changing can be done so rapidly that all meaning is lost in the work, and the rush is placed on the quantity of the output, not the quality.

Taking this example a step further, operating teams in this manner, where humans and projects are multitasking and changing tasks constantly, is counterproductive. When a product must be re-architected or a code base refactored, it’s a common trapping to forget all the newly accumulated technical debt. A tendency is to keep pushing forward, without being mindful of the train we’re pulling behind ourselves. A train can’t stop and take a 90-degree corner, and it’s not practical to pick every train car up and reseal it on the tracks when a product takes a turn. Good product owners and engineers know that it’s important to get all the work done before the train arrives. You can’t lay track under the train and you can’t survey your route one mile at a time. This paper doesn’t assert that we make less changes or stop enhancing and evolving our products, but rather that a longer view is necessary when considering the prioritization of features and functionalities that yield the best long-term outcome. All the short-term planning in the world can’t make up for a poorly managed long-term goal, or a product that isn’t ready for customers. Directional changes need to be planned and managed, especially including solutions for moving the platform and keeping it together through these transitions. Very often the focus is on the locomotive doing all the pulling up front, but if the cars fall off the tracks behind it, any first mover advantage is jeopardized. Maybe the next train can survive without derailment, but isn’t it more important for customers to get it right the first time?

## **1.3. The Eighty Percent Rule**

To create a construct to help balance progress and stability, the Eighty Percent Rule is a good tool to evaluate both micro and macro levels of decision making related to industrial innovation. It’s instructive to not get caught up “chasing the nines” when building a proof of concept, but rather to start out with a much looser framework. “Chasing the nines” as it relates to product reliability and uptime is something that comes later on down the development path of mature products. In early product stages, the Eighty

Percent Rule paves the way to make active decisions about overall readiness for providing the nines of reliability. Only after obtaining this benchmark is it appropriate to begin working towards that plateau of uptime and reliability.

The Eighty Percent Rule can be loosely defined as a means to strike a balance between efficiency and speed. The guiding idea is that there is a need for product owners to honestly rate the product experience, from the position of the consumer, and knowing its technical shortcomings. Simply put, if you wouldn't give your product an honest B rating, then it's not ready to be put in front of customers, let alone mass deployment and exposure. If a feature isn't going to function at 80% efficiency or more, it should probably hold.

The Eighty Percent Rule is also helpful for resolving discrepancies – because it's usually impractical to think in terms of getting something absolutely right. Corrections are expected as part of the process, so the aim is to get it mostly right, in order to resolve gridlock and move forward. This also serves as a nice goalpost to evaluate forward progress. The Eighty Percent Rule is designed to favor and reward long-term planning. It would be a foolish task to think we're going to show up and win a marathon when we haven't done any preparation for the race, so it's important to know where you're going in order to steer your product farther down the road. One of the tradeoffs of the Eighty Percent Rule is that it can take time to build the necessary momentum to get a product launched; it's important to keep in mind that this construct doesn't necessarily include a fast start, and it takes time to move with precision. If products are built with these key concepts, the time investment to moving quickly is minimized.

#### **1.4. Hitting the Moving Target**

The final piece of the puzzle requires our leaders to think in three dimensions. At a critical junction, a product goes from a small patch of trial sites to a multi-million-dollar production across dozens of geographic locations. At this point scale becomes a more heavily-weighted factor when making decisions about a product. The more knobs turned, the longer corrections will take to execute -- but oftentimes product leaders can inadvertently forget about reality testing decisions that can get made in haste. Was sufficient time budgeted for changes to be cascaded throughout the system? Was time budgeted to pay off technical debts, or did the trajectory lead straight into the next development hurdle? This methodology can appear to be cumbersome on the surface. In practice, however, managing low technical debt prevents a product or a team from falling into the abyss. Changes that are more deliberate and paced tend to yield better decisions about future planning. A useful goal is to consider the best decisions for the month, not the moment.

A premature commitment to a product build can have the same effect as changing a product too rapidly. Essentially, it's two sides of the same issue: Scaling an enterprise product is not assembly line work, and the best way to do that, especially with a lean team, is to keep the product flexible. When the product reaches this phase, it's useful to focus on the realities of exponential scale. It's no longer trivial to make changes to a few sites by hand anymore, or to rebuild the system from scratch. Any change is multiplied by the size of the deployment footprint. Making an architectural change at scale? That's going to be costly to go back and redo all those sites a second time. Need to make a global configuration change? Even a simple task can take hours when multiplied over a large footprint, let alone complex changes. Have to make several changes to a product? Validating that all those changes were made correctly can be error-prone, especially if a human must do all that swivel chair work. Incredible amounts of cost and resource losses can accrue very quickly at this stage -- because hitting a moving target is difficult. We have to be mindful about predicting where we're going to be upon impact. A simple miscalculation can cause a missed target -- which results in a wasted first attempt, as well as time lost to take another shot.



When time to market is critical, making the first shot will yield the best return on investment. Customers won't be impressed by a half-baked product, and first impressions are critical to new product adoption.

## **2. Launching A New Product**

### **2.1. Universal Truths**

The information outlined in subsequent sections of this paper intend to reinforce or structuralize what it takes to rapidly deploy a new product to the market in the cable sector. These tenets are reasonably ubiquitous, regardless of work portfolio, and can be considered common regardless across frameworks.

### **2.2. Challenge of Scale**

#### **2.2.1. *Inventory and Accountability***

In 2020, some take for granted that modern cloud computing products have built-in reporting, monitoring, or telemetry features. In some products, the physical layer has been completely abstracted, so that consumers don't have to manage that layer of infrastructure. But when a company sets out to build a new product at the edge, a lot of those efficiencies stop at remote access, and each hardware manufacturer has its own interface and operating system. There may be a common management interface, or SNMP trapping, but it's worth recognizing that to the operations product owners, it looks like a blank slate. If the hardware platform is an internal build, it's most likely going to require specific firmware revisions, or specific driver requirements. When this hardware platform serves dozens or thousands of devices, visibility across the footprint is necessary, to audit for discrepancies. At a minimum, telemetry should be built into the product, to inventory critical platform components, such as: installed RAM, SSD wear rate, drivers, firmware, BIOS, fan status, power supply status, and temperature alarms, among others. Such visibility is a minimum requirement -- for further efficiencies, auditing capabilities can be added to actively seek out exceptions and raise their visibility, resulting in corrective action. Mastery of cloud components includes the ability to not just audit the footprint, but also to take components out of service for patching with minimal to zero intervention. After all, once the product is running, it's likely going to be run by a lean team, so this functionality will realize a vast amount of utility through the end of the lifecycle of the product. It's important to be able to build into a product the ability, early on, to know how many exist, where they are, and their current state. In order to move faster, having the ability to orchestrate the hardware layer is an essential foundation to any enterprise product. For best results, it is one of the first components at the integration layer, built after the proof of concept is working.

#### **2.2.2. *Monitoring***

Monitoring is another component that is often overlooked until the product gets closer to launching to customers. There are advantages to beginning to monitor the hardware installations coincident with the system being built, but be careful -- this can also be taken a bit too far. Example: A product launch involving SSD drives, which, in older installations, had completely worn themselves out -- before it was ever used for customers. Had this been monitored earlier, the issue could have been identified with an immediate impact assessment because of good inventory practices. This kind of surprise usually happens near the moment when the equipment is lit up with customers, and it throws another technical barrier in front of the product launch. And because the solution usually involves coordinating either with a vendor on site or asking for help from local site personnel, it's not exactly the fastest procedure, and usually ends up involving multiple resources. At this point there are usually a host of other critical issues being

addressed before launching the product -- a good project lead will work to identify the platform's critical components, and build systems to monitor and assess system health so that a lean team can operate the platform with maximum efficiency.

While it is critical to build monitoring early and often, it is also helpful to suppress or mute production alarms while the system is being built. If the system is in a full production monitoring mode before customers are on the equipment, that usually means change management tickets need to be coordinated to make updates to the platform. This creates additional overhead by means of adding additional layers of process to the workflow, and across many sites this can add up quickly. While the system is greenfield (in this case, meaning no customers are on it), the platform may need to change frequently to keep up with the pace of the developers. For best results, build alarming into the system early and schedule regular reviews of the alarms to make sure the system is behaving as expected. But don't paint your teams into a corner by creating additional administrative overhead for anyone making changes to the platform when there is no exposure to customers. A good monitoring system, when implemented early, can help identify problems in the platform and surface them before they become critical. Maintaining a platform, at scale, while building it, carries a higher up-front cost, but the payoff resembles that data center experience that every product wants to operate in. By conquering the hardware footprint with good inventory management, and having the ability to automate upgrades and by monitoring those investments, products are in a very good position to sustain a much more rapid CI/CD model.

## **2.3. Challenge of Change**

### **2.3.1. *Data In, Data Out***

This paper will conclude with some insight about users and user experiences, as it relates to expediting customer hardware deployments. At this stage, it's useful to keep in mind that it is when products scale that the most can be learned about the platforms that were designed. It's unacceptable to just have the data if it can't also be managed. CI/CD wants to always be moving. In order to stay highly effective, which is to say moving without stumbling, any management systems that are built should prioritize consideration of the teams who will ultimately run or use the platform, day after day after day. Ideally, both back end and front end teams coordinate on the architecture and data flows, to make the language and technology decisions. After that, how do the teams go about building the system and adding components? Do they build a UI that allows for manual data entry, so configurations can be managed in a central database? Or do they build a protocol and an interface to onboard devices to an endpoint, and have devices send their configurations into a central repository? The best answers here depend on how the data needs to be maintained, and there may not be a wrong answer. Now that this large data set has been created, do the teams have the ability to make bulk updates across the platform? Not everything needs to be mutable, but building it such that the most common data points that can change are scalable is useful. Also: Does the system expect that it can be handled by a database admin, or are end users also empowered to manage the data?

Building on the previous tenets of inventory and monitoring, data is the next piece of the system that must be reinforced. This user interface may account for the needs of users for several years to come, so this step is critical to getting right. The cold reality is that once a product is launched and there are customers on it, operators tend to become risk averse. This can sometimes mean having to settle for a mediocre UI. A good timeline to keep in mind is about ten years. If the platform can be put to good use for about a decade, before having to rebuild it, that's a good system. Such timelines are also useful for future planning: can any shortcomings of the system be tolerated for the next ten years? Or will running the platform always require more resourcing? How does one not only support the platform, but any syndication partners, too? What kind of example does this set within a company if its enterprise tools are

subjected to compromise and work-arounds, just to perform daily functions? Put another way, does the platform inspire your teams to do their best work?

### **2.3.2. *The Hardware Lifecycle***

Nothing in the hardware computing world lasts forever. In 2020, a best-case scenario for a syndicated product is a lifespan of about a ten years. Some older systems in the industry, built in the early 2000s, could have easily been designed to last ten years -- but as the product space for today's technology gets more crowded, stagnation becomes unprofitable. Old equipment and legacy technologies will end up costing money in the long run. Because the pace of innovation and competition is relentlessly increasing, hardware platforms are unlikely to have the same lifespan they enjoyed twenty years ago.

It may not be possible to know what a replacement hardware platform will be, but when building a platform, it's useful to ask suppliers questions about how long that model of equipment will be in service. There are no guarantees, and it can happen that a custom hardware product may be use equipment that is near the end of the manufacturer's hardware refresh cycle. Or, the opposite can happen, in that the appliance could enjoy a foreseeably favorable lifecycle. Gathering data points that can inform an estimate about your hardware's expected life span will help guide any decision making about how much time and energy are invested into the platform. While we don't want to withhold functionality or usability from the users, it may be prudent to investigate what compromises should be made if the platform is going to have to be rebuilt for technical reasons and in short order.

### **2.4. Meltdown Intel Security Vulnerability – Retrospective**

Catastrophic events in the digital world are uncommon, but they aren't impossible. Looking at a fairly recent example, the Meltdown Intel security vulnerability, discovered in 2018, triggered major impacts on the hardware and software world – especially for delivering video. When this vulnerability was disclosed, it brought home a very real problem to an internal project involving thousands of Intel Xenon CPU processors, distributed across dozens of locations, to do real time linear encryption and ad insertion. The performance of the product, at the time, was calculated based on a a certain density of the hardware solution which met the product's architectural needs for streaming throughput. After the Meltdown vulnerability was published, the resulting software patching of this vulnerability generally resulted in a performance loss of the CPU. If your product is very CPU dependent, a performance loss can have major impacts on its ability to service the needs of your customers. These kinds of incidents can have major architectural impacts to custom hardware products. This kind of vulnerability doesn't happen often, but that such a sizeable one did within the last five years is a good reminder about diligence in system design.

If you have a CPU intensive product, and something like Meltdown happens again, how do you quickly evaluate what options are available? How does your product recover and move forward? What if it costs money because in choosing not to compromise on security, you trade off being able to support the service levels of syndication contracts? Is it possible or cost effective to throw more CPU at the problem to make up for the shortfall? Do you focus on gaining efficiencies from the software and try to get that to be more performant? Or does the product team need to look at other options, to determine whether there is a critical tipping point that will need to be crossed in the design? As the investment into a product increases, it behooves product owners to be able to immediately materialize specific information about their platform to make the most informed decisions possible. If another Meltdown happens tomorrow, does the product have the telemetry in place to support good decision making? With so much invested into building these products, it's borderline negligent to build products without telemetry. Product owners should be able to assemble a data set of information about their platform in time for their next meeting, not in a few days. This same principal applies to other consumables for your product -- whether

it's CPU cycles or network bandwidth, it's going to be critical to have real data about your product to make the best decision to build it on or ahead of schedule.

## **2.5. Who Can Move My Cheese?**

Each product is different, and with this uniqueness comes a host of common and individual characteristics that are important to manage for any given system integration. The previous tenets of this paper discussed how essential inventory and monitoring are to building a good product. If there is an outage and teams are escalated to troubleshoot a problem, time is money, and finding the problem becomes a race. A well-built product, that has visibility into all the code bases across the footprint, can quickly help rule out problems and steer troubleshooting in a positive direction. If the product can detect that there's a firmware mismatch, or that a node failed an upgrade, this can lead the production teams to the source much more efficiently. On the other hand, if there is a good auditing system and all the basic health checks for the system are coming back clean, then we can assume the problem lies elsewhere and quickly rule out the first layer(s) of the product. When systems grow large and there are thousands of data points to look at, humans are going to be not well equipped to spot those small exceptions.

On top of the physical layer, enterprise products are typically built with redundancy in mind. Oftentimes modern products are geo-redundant (meaning there are multiple co-locations where the loss of a physical facility can be carried by another peer). Is it possible to build this logical layer into the management of the product, so it can understand product states? Can the platform be upgraded without requiring a skilled operator who understands the specifics of failing the systems over, in order to perform work on the platform? The reminder here is that operations teams are generally lean, so anything we can do as product builders to make common tasks easier, or find a way to decrease the amount of time and skill required to accomplish maintenance work, is beneficial. A core theme for this paper is that to encourage and promote CI/CD, we **MUST** be designing products platforms that are easy to upgrade and efficient to operate. We need to eliminate the technical debt so efficiency is baked into the system. There's a benefit to being able to build and scale a platform with small teams, so the small teams of operators are best equipped for this task. Time and time again in cable we make the same mistakes by pushing unsupported products to market too quickly. We fail to recognize that the lean teams handling the build, the customer transition, and the operation of the platform are not going to be the best suited to find time and necessary skills to work backwards through the hardware stack to create a useful orchestration layer, after a product has been built. This is the crux of the issue -- products are created in a vacuum, then pushed out the door without proper supportive tooling. They change too frequently before getting in front of customers, and lack a good orchestration layer to manage these changes. Products end up behind schedule because they aren't built in a way that helps operators get them to customers faster. We must collectively stop this behavior if we expect to improve our time to market. We must stop believing that as leads, our job is done once the proof of concept is working in the lab. We must be invested in the management of the platform and product that comes after this idea.

## **3. System Users**

### **3.1. Someone Else's Problem**

Traditional product lifecycles dedicate much of their development to creating a working product, and frequently this comes at the expense of the end users. As developers, we often lose sight of the fact that the product isn't complete until we have system users and managers. The R&D road to successfully solving a new problem is often a challenging journey, and it's a common trapping to collapse across the

wrong finish line. So much time is focused on the problem that we forget to also deliver a solution that works for the end users. Engineers can become hyper-focused with building a working proof of concept and are allowed to completely ignore the complexity of the system they've built. Very often, the mindset of being aware or caring about the next phases of a project differs from the traditional build mindset. Often the product builders don't have any connection to the UI, the users of the system, or the operations of the system, which become a problem for "someone else". Or, more formal barriers emerge and takes the incarnation of a handoff between product development teams and operations teams, further exacerbating the problem. Building a management and orchestration layer then becomes automatically inherited technical debt. Operations teams aren't typically staffed or funded to build a UI that meets the complex needs of its users. Each of us in the industry, not to mention our end customers, can probably think of a product we worked on which had a compromised UI. The product itself works just fine but using the system to manage the product can come with a lot of work arounds or hang ups. This isn't the hallmark of an enterprise or world-class product, and as project leads aiming to improve our speed to market, it will be imperative to think through these challenges in a new way. The team who built the prototype typically has the most working experience building the product. Why wouldn't we leverage this experience and have this team construct the UI or the basic foundations for manipulating data within the system? Instead we allow these experienced teams to discard this knowledge and then the snowflake becomes a snowball rolling down the hill. The last barrier for innovation and correction of a management platform is exposing the product to live customers. The appetite for product owners to make big changes to their platform once customers are stably onboard approaches zero in a short time span. The risk vs reward construct comes into play, and with paying customers on your product, that is the most prudent business decision. We then find ourselves in a place where customer migrations override any needs for the operations teams to manage the product, and the window for making changes to the platform closes like a feedback loop, underscoring the importance of good tooling.

### **3.2. Accessing, Viewing, and Changing Data**

Big changes don't typically happen to a product after it has been launched to customers, but in the beginning stages of a product, large changes and disruptions are much more common to the landscape. Couple this with productivity methodologies like CI/CD, which aim to continually keep making changes, and we have a recipe for creating a lot of technical debt. Often, we find that our engineering assumptions don't always happen in the real world, and we're forced to make pivots to designs and plans when new information presents itself. This is a common theme for prototyping – we make enough assumptions to move forward and reserve resources for when that doesn't always work out. When we must make these corrections, how will the middle and end users be able to access, view, and manipulate the product's data? At scale, changing data by hand is inefficient and impractical, so we need methods of importing and exporting data if they can't be manipulated in the tool directly. How will other users or teams consume the data your product will generate? Will other users or systems have access to the data? It's imperative to be able to make scale changes to a product, especially in the early stages of its deployment. Product leads need to be finding ways to minimize technical debt from the earliest stages of a product's lifecycle in order to sustain the changes needed to mature the product.

Balancing the needs of the product's users and the development cycles required to achieve those goals requires a delicate balance to be struck- products can't wait forever before they need to be exposed to real world challenges. With some foresight and experience, we can anticipate what our user's needs will be and find compromises in functionality that help us achieve rapid product penetration.

### **3.3. Why Did We Wait This Long?**

Waiting to add users in to the closing stages of your product's development is an open invitation for rework, and by nature of what's been discussed in this paper, a challenging exercise with diminishing opportunities and resources. There is logic to building a prototype first before creating a UI, but we must acknowledge that there is a lot of development time ahead that need to be spent transforming a prototype into a world-class product that's not only fit for your customers, but valuable enough to market to syndication partners. There is a flexible window in building a product to introduce your users and start building for their needs. But when prototyping new products, we should be doing the basic research to anticipate what the users will need, and what data is going to be valuable, ahead of time, then bake that into the fabric of the product. We need to be building platforms that have the full vision for what the product needs to do, how it should behave, and how it's going to be used, from the earliest stages of conception. The users are what bring the product alive -- without them, it's just a bunch of ones and zeros.

### **3.4. Build a User Interface That is Meaningful; Build Your UI in a Meaningful Way**

When you interact with something that is designed well, it feels natural and intuitive for how to accomplish a given task. UIs that don't feel obvious leave users frustrated and can create barriers to adoption. It's important to represent your company's investment by creating a UI that works well and works well for your users. A UI will seldom be perfect, as that is an objective measure, but having a UI that is in agreement with what your users want and need is central to the long-term success of the product. A UI represents your product, and a misrepresented UI can leave a user speculating about the effectiveness of the underlying technology. It's the paint job on the car, or the finishing materials selected in the interior. A sloppy build on the surface will be an implied reflection on the underlying technology. The UI should do all the talking and speak for itself, leaving all explanations absent.

As important it is to have a natural interface, it's also crucial to make sure that the UI can do all the heavy lifting for your users. If large, bulk updates must be made to the platform in order to maintain it, the UI needs to have a way to orchestrate these large changes without needing to involve a back-end user to facilitate the change. End users must be empowered to care for and manage their platform. If doing so is hard to use and time consuming, the userbase will not be inclined to maintain non-critical data. Working on the product will become tedious, and the employees supporting it will only do the minimum required to "keep the lights on."

## **4. Conclusion**

Whatever productivity suite we're operating in for the moment is only as effective as the experience it's being built on top of. If we lack the vision and experience of working on a product through all stages of its life, we can compartmentalize deficits and defer responsibility without consequences. All this comes at a cost, and these assumptions we make about other teams who will pick up the slack usually means that there's not a coherent plan or vision for the lifecycle of the product. If we have no common vision or goal, and no means to get there quickly, we must expect that projects will meander through time and will eventually arrive in front of our customers. If we want to empower ourselves to take control of these timelines, we must start by being honest about planning around the full spectrum of the product's lifecycle at every stage with every team involved. We can't unfurl the Mission Accomplished banner across the windshield just because we solved the prototype challenge -- there's a long road ahead to taking this fledgling product to millions of customers. We find that when projects are run well and have this vision in mind from the beginning, they often start slowly, but finish quickly. It takes time to invest

in orchestration, management, and user requirements, and when platforms are being built, they are often done so through lots of learning (see: failure) and chaotic change. Product leads need to be mindful to create and build orchestration and management into the platform so that navigating all these changes and managing technical debt is factored into the platform. Our modern productivity suites emphasize rapid and dynamic changes, and without a proper support system for this disruption, we can quickly find our products in the middle of a vortex.

Building a product should be done so with the least technical and experienced users in mind. This is a great pipeline for engineering talent -- more experienced engineers design and build a platform, and more junior engineers can learn how to operate and support the platform. It's a great way to get to know your product and business from the inside out. As these engineers grow, they accumulate experience and knowledge of how a good product is built and operated, thus helping to create the next crop of engineers and leaders in the company. To do so, we must create meaningful UIs that make the operators passionate about their product and give them a sense of authority and accomplishment. We can't expect every end user to be a database administrator, therefore we need to build our systems for the most common denominator and prioritize functionality that will empower our engineers to be a force multiplier, especially considering that lean teams are typically "keeping the lights on" for these products.

By nature, prototyping a new product means to boldly go where no person has gone before, and building new world class products doesn't happen overnight. Often, it takes years of work to find the balance between budgets, expectations, technology, and resources to craft a meaningful middle layer to manage your architecture. Being able to keep an enterprise platform running at scale with a lean team demands that all the hard work be done up front and most of the challenges and desires of the users can be met autonomously. Platforms are born, and they are born to change. Flexibility, inventory, and scale are critical components to master when building and finishing building your platform. Hardware is going to change due to natural or unnatural causes; additional technology disruptors like those that surfaced in recent history are practically inevitable. The most effective product leads know how critical it is to account for consumables in real time and are robust enough to support large changes across the footprint in a short period of time.

In conclusion, project leads shouldn't wait too long to invite users into your system to find out all the ways your they will need to access and consume data from the platform. Establish relationships early and engage users directly about how they think the system should behave and what their pain points are with the existing system. For best results, build a user interface that removes technical burden from the operations teams so they can be left to an already important task of running, monitoring, troubleshooting, and repairing their platform twenty-four hours per day, 365 days per year. Product leads need to be mindful about not leaning in so far to the Continuous Delivery that we forget that the other half of the feedback loop is Continuous Integration. We can't simply ignore the process for the sake of meeting deadlines and pushing an unsatisfactory product out to our customers too early. Product leads need to incorporate tooling from the ground up and across team handoffs in order to build a product that can get to market the fastest and most direct route possible.

## Abbreviations

CDN	Content Delivery Network
CI/CD	Continuous Integration / Continuous Delivery
CPU	Central Processing Unit

MSO	Multiple Systems Operator
TCP	Transmission Control Protocol
UI	User Interface

## Bibliography & References

CI/CD Web Article     <https://www.infoworld.com/article/3271126/what-is-cicd-continuous-integration-and-continuous-delivery-explained.html>

Cloud Market Share 2020     <https://kinsta.com/blog/cloud-market-share/>

Meltdown and Spectre     <https://meltdownattack.com/>

Still Fighting Meltdown and Spectre     <https://www.wired.com/story/intel-meltdown-spectre-storm/>

Shane Pape, Mgr 2, Prodt Dev Engineering, Comcast. Interview.



# **RF Testing Applications for Software-Defined Radio**

## **Using SDR for Impairment Generation**

A Technical Paper prepared for SCTE•ISBE by

**Robert M. Lund**  
Principal Engineer  
Next Generation Access Networks  
Comcast  
720-512-3691  
Robert\_Lund@comcast.com

**Kathryn Sanders**  
Intern / UIUC Undergraduate Student  
Sanders RF Consulting LLC  
Kgs5@illinois.edu

# Table of Contents

Title	Page Number
1. Introduction .....	3
2. Software-Defined Radio General Architecture .....	3
3. SDR Software: GNU Radio .....	5
4. Sampling Theory .....	6
5. The Discrete Fourier Transform or DFT .....	8
6. Introduction to Profile Management Application .....	10
7. Practical Application Example #1– Capturing LTE Signals For Replay as an Impairment Source .....	12
8. Practical Application Example #2 – How SDR Has Aided Profile Management Application (PMA) Testing.....	15
9. Equipment Used .....	22
10. Conclusion .....	23
Abbreviations.....	24
Bibliography & References .....	25

## List of Figures

Title	Page Number
Figure 1 – SDR Architecture Block Diagram.....	4
Figure 2 - GNU Radio Flowgraph Example .....	6
Figure 3 – Aliasing Due to Undersampling .....	7
Figure 4 – Square Wave Construction From Sine Waves.....	8
Figure 5 – The Profile Management Application System .....	11
Figure 6 – PMA In Action: Building Custom OFDM Profiles .....	11
Figure 7 - Finding LTE Band In Use on iPhone .....	12
Figure 8 - Real Time Spectrum Analyzers for IQ Capture (clockwise): osmocom_fft, uhd_fft, gqrx SDR .....	14
Figure 9 - Combining of Noise Sources in Laboratory Environment .....	14
Figure 10 - GRC Flowgraph to Construct and Transmit Custom Impairment Waveform.....	17
Figure 11 - Graph of RxMER Data Generated by Embedded Python Block.....	18
Figure 12 - Qt GUI Frequency Sink Visualization of Impairment Signal.....	19
Figure 13 - Qt Fospor Sink Visualization of Impairment Signal .....	19
Figure 14 - Bandwidth Monitor of IQ Data Stream.....	20
Figure 15 - Impairment Signal on Spectrum Analyzer.....	21
Figure 16 - OFDM Profile for SDR Generated Impairment Signal .....	21

## List of Tables

Title	Page Number
Table 1 – LTE Frequency Bands (3) .....	13
Table 2 – RxMER TFTP File Format (4).....	16

## 1. Introduction

Recent advances in digital communications algorithms and artificial intelligence, coupled with the exponential growth of personal computing power, have positioned software-defined radio (SDR) technology as a viable tool for RF signal capture and playback, simulation, and testing. Popular applications for modern SDR designs include:

- Cellular handsets with programmable digital signal processing (DSP)
- University research / teaching aids
- Military and satellite communications platforms that make use of programmable cores for intermediate frequency (IF) and baseband signal processing
- Cognitive radio – a “smart” radio mesh that can adapt to interference adjusting frequency or power for example
- Amateur radio

The goal of this paper is to outline the general architecture and capabilities of a software-defined radio, followed by some basic principles of digital signal processing. A familiarity with sampling theory and the discrete Fourier transform, for example, will be very helpful for anyone looking to get started with SDR.

Finally, two lab applications of software-defined radio will be detailed. Both of these applications were used in testing a profile management application implementation, so a section on PMA is introduced. After familiarization with PMA, the first test example uses SDR to capture long term evolution (LTE) signals, then re-play them back into the RF plant to be subsequently detected by pattern recognition software that is a function of PMA. The second case recreates plant conditions from a cable modem’s perspective by translating its reported receive modulation error ratio (RxMER) values into a waveform to be used as an impairment profile applied to an orthogonal frequency division multiplexing (OFDM) channel. Again, this impairment can then be analyzed by the analytics engine of PMA, and a custom OFDM profile recommendation can be verified.

## 2. Software-Defined Radio General Architecture

A software-defined radio is a radio that replaces traditional hardware components such as mixers, filters, and modulator/demodulators with software implementations, typically running on personal computers and driving a minimalistic hardware set. The Wireless Innovation Forum and IEEE P1900.1 working group has further simplified the definition down to:

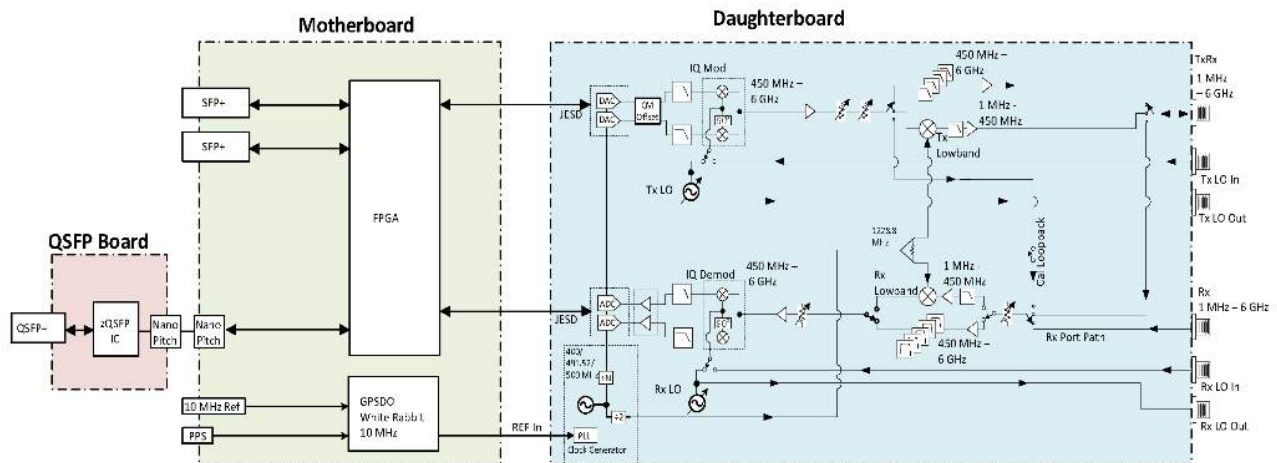
*"Radio in which some or all of the physical layer functions are software defined" (1)*

There are many advantages to software-defined radio designs. Potentially first among these is the low cost. Devices can range from \$10 USB-powered dongles, up to \$10,000 high performance units with multiple transmit/receive (TX/RX) channels, higher maximum sampling rates, and greater usable bandwidth. Another primary advantage is flexibility. SDRs are not hardware limited to any particular specification standard, modulation type, or operating frequency range,

for example. The ability to test new designs in software also aids in the rapid prototyping of communication systems.

Conceptually, an ideal software-defined radio would eliminate any analog signal processing stages, aside from an antenna, power amplifier, and input/output (I/O) sink or source such as a speaker or microphone. In reality, today's SDRs use an analog frequency conversion stage at the receiver, with all signal processing thereafter being digital. Similarly at the transmitter, the digital-to-analog conversion stage is followed by local oscillator/mixer for transmit frequency conversion along with power amplification.

To illustrate a typical SDR design, the block diagram in Figure 1 outlines the functional components. This was the hardware used in the two practical example sections to follow.



**Figure 1 – SDR Architecture Block Diagram**

The key features of this design include:

- Xilinx Zynq-7100 FPGA SoC
- Dual-core ARM A9 800 MHz CPU
- Two RX, two TX in half-wide RU form factor
- 3 MHz to 6 GHz frequency range
- Up to 200 MHz of instantaneous bandwidth per channel
- Sub-octave RX, TX filter bank
- 14 bit ADC, 16 bit DAC
- Configurable sample rates: 200, 245.76, 250 MS/s
- Two SFP+ ports (1 GbE, 10 GbE, Aurora, White Rabbit)
- One QSFP+ port ( 2x 10 Gb / Aurora )
- RJ45 (1 GbE)
- 10 MHz clock reference
- PPS time reference

- External RX, TX LO input ports
- Built-in GPSDO

One of the main factors to consider in selecting a radio is usable bandwidth. The specification above lists 200 MHz of available bandwidth that will allow us to simulate or impair a full OFDM block with 190 MHz of modulated spectrum. Another point to consider is sampling rate. Radios range from fixed, to selectable, to fully configurable sampling rates, with the fully configurable ones being the most flexible, because in some cases they can eliminate the need for additional resampling functions. Many designs use open-source Verilog code for the FPGA that can be modified for high performance applications.

### 3. SDR Software: GNU Radio

GNU Radio is a free open-source development project that implements signal processing blocks in software for waveform creation, sampling, and analysis, that supports a variety of software-defined radio hardware. It is used extensively in academic and commercial environments for wireless communications research and real-world radio systems.


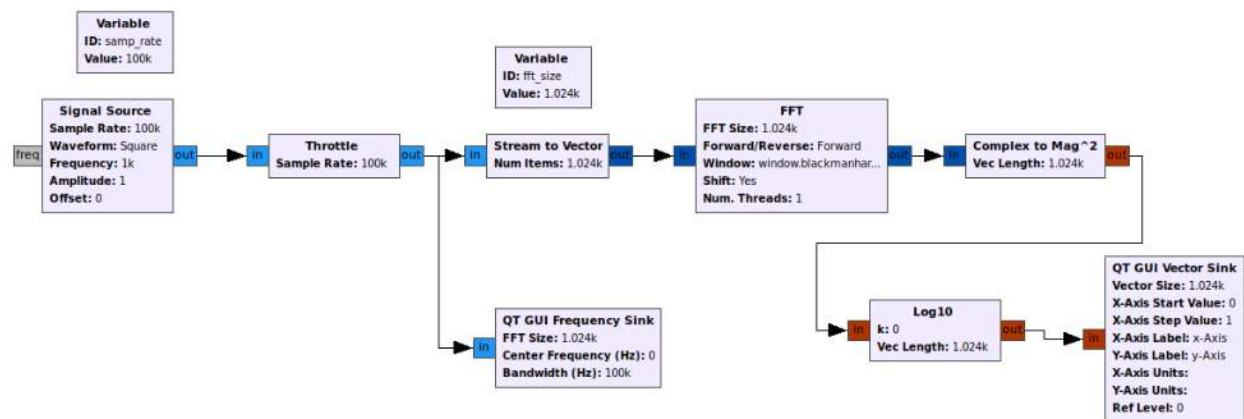
GNU Radio consists of software-based signal processing blocks used in conjunction with hardware drivers to build applications capable of transmitting or receiving digital data streams. These streams can be directed to disk and saved in different file formats, or to a software-defined radio sink. GNU Radio resource blocks include equalizers, filters, modulators/demodulators, filters, and other common radio elements. These blocks' inputs and outputs are then connected in software to build an end-to-end flow.

Being software based, GNU Radio operates on digital data. It uses complex baseband samples as input for receivers and as the output data type for transmitters. Analog mixing hardware can then shift the signal to the proper frequency.

One of the biggest advantages to working with GNU Radio is that it is primarily written using the Python programming language, with the computationally intensive operations written in C++. GNU Radio also incorporates a graphical user interface – creatively named

GNU Radio Companion – that offers drag-and-drop software blocks to construct flowgraph applications. These GRC flowgraphs present an intuitive approach to thinking about the discrete signal processing stages. Once a flowgraph contains a source and sink, the file can be saved and executed, and Python code is automatically generated. At this point there is no need to run the flow from within GNU Radio Companion, as the Python code can be executed directly from the command line or Python shell.

The GRC flowgraph shown in Figure 2 is an example of how a fast Fourier transform (FFT) block can be used to reproduce the same output as another GRC provided block: the Qt GUI frequency sink. The only difference is that Qt GUI frequency sink has the FFT, stream-to-vector, magnitude, and log functions built in.



**Figure 2 - GNU Radio Flowgraph Example**

In my experience, it is best to compile GNU Radio, SDR drivers, and associated software from source code for your specific operating system. GNU Radio can compile on Windows, Mac OS, and Linux. Ubuntu 20.04 has been used for the testing contained herein.


## 4. Sampling Theory

When beginning to work with software-defined radio, it may be helpful to review some basic principles of signal processing.

In digital signal processing, sampling of a continuous-time input signal is performed in an analog-to-digital converter (ADC). The output of the converter is a series of data points that sample the instantaneous value of the input signal, and are taken at a given sampling rate. These discrete samples may be used to perform filtering or other processing in the digital domain using a digital signal processor, and then converted back to the analog domain using a digital to analog converter (DAC) as needed. Key to the sampling process is the sampling rate, or the number of samples taken per second (measured in Hz).

Different sampling rates are specified for different purposes. For example, 44.1 kHz is the standard sampling rate used in CDs, and 48 kHz is the recommended sampling rate for many other applications, including audio in digital video recordings. If the signal is sampled properly, the original analog signal may be reconstructed perfectly from the digital samples.

In order to choose the sampling rate for a specific application, the concept of aliasing is important to consider. Aliasing refers to the phenomenon that results in an imperfectly reconstructed analog signal from digital samples created when the sampling rate is too low. As we can see in the illustration in Figure 3, the blue signal has been reconstructed erroneously from the red signal due to the inadequate number of samples taken.



**Figure 3 – Aliasing Due to Undersampling**

In order to avoid aliasing and accurately reconstruct a signal from its samples, the Nyquist condition must be satisfied. This condition requires that the sampling frequency  $f_s$  must be greater than or equal to two times the frequency bandwidth  $B$  of the input signal  $f(t)$ :

$$N = f_s \geq 2B$$

As long as the Nyquist criterion is met, it is theoretically possible to exactly reconstruct  $f(t)$  from its discrete samples  $f_n$ , where

$$f_n \equiv f(nT), -\infty < n < \infty$$

The samples of  $f_n$  are spaced at time  $t = nT$ , where  $T$  is the sampling interval  $\frac{1}{f_s}$ .

Under the Nyquist sampling condition,  $f(t)$  can be reconstructed from its discrete samples  $f_n$  by using the reconstruction formula:

$$f(t) = \sum_n f_n \text{sinc}\left(\frac{\pi}{T}(t - nT)\right)$$

The sinc function,  $\text{sinc}(t) = \frac{\sin(t)}{t}$  is also known as the interpolating function.

## 5. The Discrete Fourier Transform or DFT

Mathematically, the Fourier series is used to represent any periodic function as a sum of sine and cosine functions. If we have any  $f(t)$  with period  $T = \frac{2\pi}{\omega_0}$ , then the Fourier series is

$$f(t) = \sum_{n=-\infty}^{\infty} F_n e^{jn\omega_0 t}$$

with the Fourier coefficients

$$F_n = \frac{1}{T} \int_T f(t) e^{-jn\omega_0 t} dt$$


where Euler's formula is used to give us our sinusoids:

$$e^{jn\omega_0 t} = \cos(n\omega_0 t) + j\sin(n\omega_0 t)$$

where  $\omega_0$  represents frequency,  $t$  represents time, and  $j$  represents the imaginary number  $j = \sqrt{-1}$ .

To visualize the construction of a periodic function as a sum of sines, the animation in Figure 4 depicts the construction of a square wave that becomes closer to the ideal with higher frequency resolution:

**Figure 4 – Square Wave Construction From Sine Waves**





If a function is aperiodic, it can still be represented as a sum of sinusoids, but the frequencies of the sinusoids are no longer simply harmonics of a fundamental  $\omega_0$ . A general aperiodic signal consists of a continuous spectrum of frequencies, so rather than integrating over a single period of  $f(t)$  as is the case with Fourier series coefficients, we must integrate  $f(t)$  over all time to derive the exact Fourier transform, which is the equivalent to the Fourier series coefficients but

$$T = \frac{2\pi}{\omega_0}$$

in a continuous frequency range. In the limit as the period  $\omega_0$  approaches infinity in an aperiodic signal, then the fundamental frequency  $\omega_0$  approaches 0, and so we must integrate over the entire period of an aperiodic signal, which in other words, means we integrate over all time:

$$F(\omega) = \int_{-\infty}^{\infty} f(t)e^{-j\omega t} dt, \omega \in (-\infty, \infty)$$

In analog signal processing, the Fourier transform  $F(\omega)$  is used to determine the characteristics of the signal  $f(t)$  in the frequency domain. We can determine the frequencies or spectral content present in  $f(t)$  using the Fourier transform, and manipulate or filter the signal as needed in the frequency domain.

Then, we can use an inverse Fourier transform to go back to the original time domain function from the frequency domain:

$$f(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} F(\omega)e^{j\omega t} d\omega$$

Similarly, in digital signal processing, the discrete-time Fourier transform (DTFT) is used in order to convert a continuous discrete-time signal (input sequence) into its counterpart in the continuous frequency domain. It is modeled by the equation

$$X(\omega) = \sum_{n=-\infty}^{\infty} x[n]e^{-j\omega n}, \omega \in [-\pi, \pi)$$

where  $x[n]$  represents a discrete-time input signal, and  $n$  is an integer.

The inverse DTFT is given by

$$x[n] = \frac{1}{2\pi} \int_{-\pi}^{\pi} X(\omega)e^{j\omega n} d\omega, \omega \in [-\pi, \pi)$$

Finally, the discrete Fourier transform (DFT) is simply a “sampling” of the DTFT in the frequency domain – a finite version of the DTFT taken over a specified interval or number of samples. It is modeled by the equation

$$X[k] = \sum_{n=0}^{N-1} x[n] e^{-j\frac{2\pi kn}{N}}, k = 0, 1, \dots, N - 1$$

where N represents the total number of samples.

And we can also define the inverse DFT by:

$$x[n] = \frac{1}{N} \sum_{k=0}^{N-1} X[k] e^{j\frac{2\pi kn}{N}}, n = 0, 1, \dots, N - 1$$

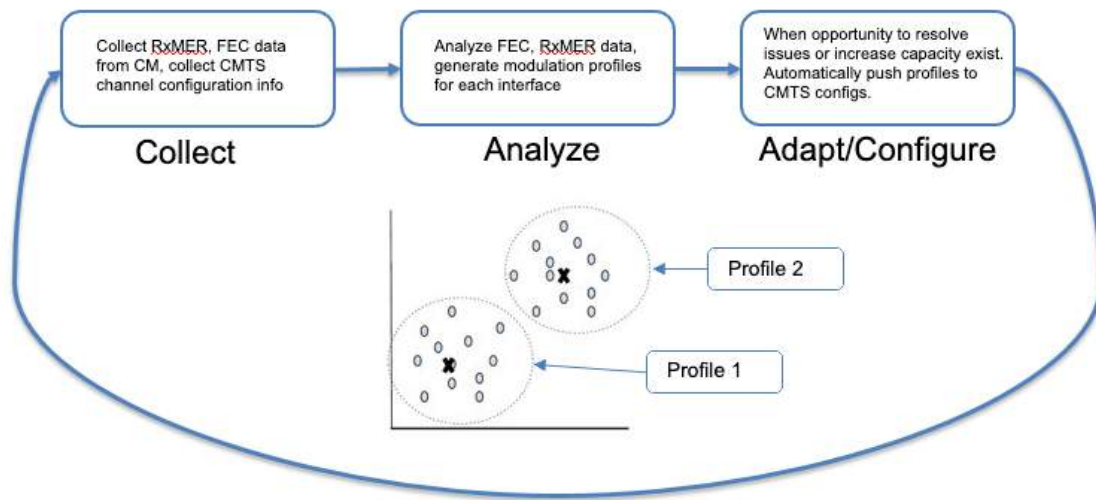
The DFT is a mainstay in the digital signal processing toolbox. It allows us to determine the spectral components of a sampled signal and to determine the frequency response of a system given the impulse response of the system in a completely digital manner. It uses a finite number of samples which can be easily used in computations of all kinds.

The FFT, or fast Fourier transform, is an algorithm that is used to calculate the DFT in a mathematically efficient manner. The FFT is employed by mathematicians, engineers, musicians, and scientists to rapidly analyze, modify, or utilize sampled signals in the frequency domain for myriad applications.

## 6. Introduction to Profile Management Application


The next two sections of this paper provide examples of using SDR to generate impairments for profile management application testing. A brief overview of PMA will demonstrate how highly customizable impairment profiles can aid in OFDM profile management testing.

A profile management application system builds an optimized set of OFDM channel modulation profiles based on current plant conditions. It does this by collecting receive modulation error ratio data from cable modems, analyzing this data to sort modems into impairment groups, and finally building the modulation profiles and applying them on the CMTS (see Figure 5). The modulation profiles themselves typically contain variable bit-loaded regions ranging from excluded regions (0-QAM) for highly impaired subcarriers, to 4096-QAM in the low noise areas of the OFDM channel. In this way, the PMA algorithm is able to maximize overall channel capacity while maintaining robust profiles that are noise-tolerant.



**Figure 5 – The Profile Management Application System**

Laboratory RF environments are often very clean, with short cable runs, limited ingress, and few taps or active elements such as amplifiers or line extenders. To test profile recommendations generated by the PMA algorithm (Figure 6) we will need to introduce impairments with variable power levels, widths, and durations. The following sections demonstrate how software-defined radio can be used to create these impairments.



**Figure 6 – PMA In Action: Building Custom OFDM Profiles**

## 7. Practical Application Example #1– Capturing LTE Signals For Replay as an Impairment Source

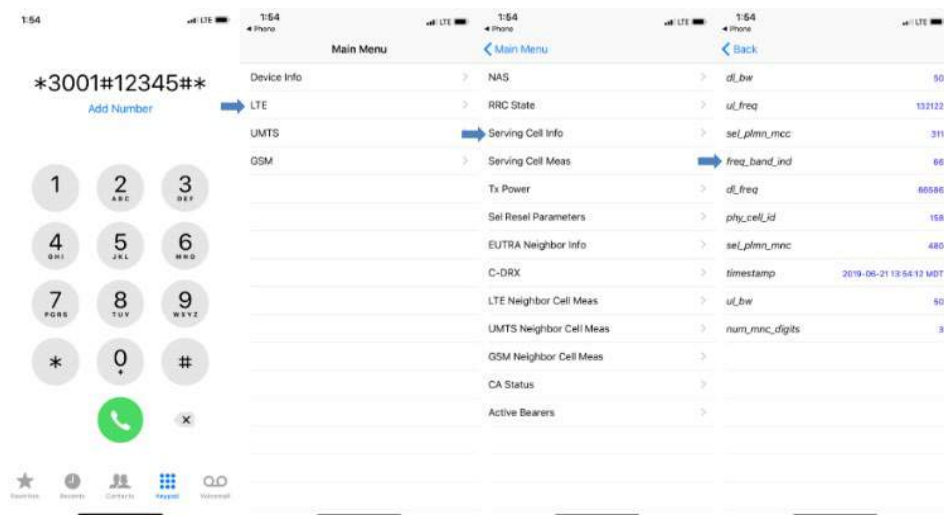
Problem Statement: How can we capture LTE signals and introduce them as ingress interference in a laboratory environment?

### SDR Solution Outline:

1. Using a cellular telephone as a signal source, find the LTE band currently in use
2. Prepare the SDR to capture at the LTE uplink frequency
3. Start a speed test on the UE to generate traffic
4. Observe/capture the LTE signal on the SDR
5. Combine SDR TX into the lab RF plant and playback capture file
6. Observe effect in the cable modem's downstream OFDM channel and verify PMA OFDM channel profile recommendation

### Detailed Steps:

1. Most modern UE devices include hidden service menu interfaces that can be accessed by dialing certain sequences. For example:
  - a. Apple iPhone (all): \*3001#12345#\*
  - b. Samsung (Android): \*##\*197328640##\* or \*#0011#
  - c. Sony (Android): \*##\*386##\* or \*##\*585\*0000##\*
  - d. HTC (Android): \*##\*7262626##\*




**Figure 7 - Finding LTE Band In Use on iPhone**

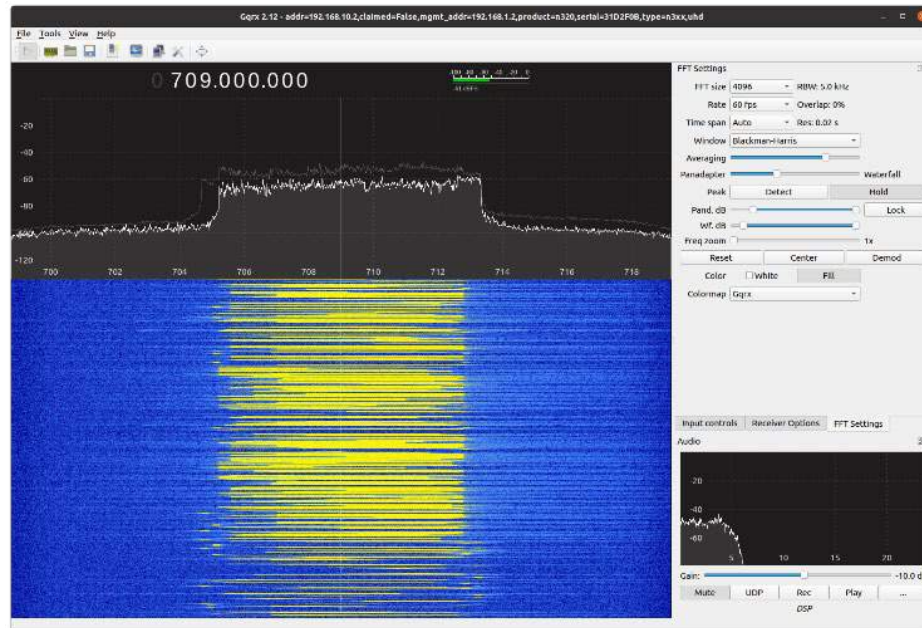
From the above example, we can determine what uplink frequencies are assigned to Band 66:

**Table 1 – LTE Frequency Bands (3)**

Band	Duplex mode[A.3]	f (MHz)	Common name	Subset of band	Uplink[A.2] (MHz)	Downlink[A.3] (MHz)	Duplex spacing (MHz)	Channel bandwidths (MHz)
2	FDD	1900	PCS[A.4]	25	1850 – 1910	1930 – 1990	80	1.4, 3, 5, 10, 15, 20
13	FDD	700	Upper SMH[A.7]		777 – 787	746 – 756	–31	5, 10
66	FDD	1700	Extended AWS (AWS-1–3)[A.17]		1710 – 1780	2110 – 2200[2]	400	1.4, 3, 5, 10, 15, 20

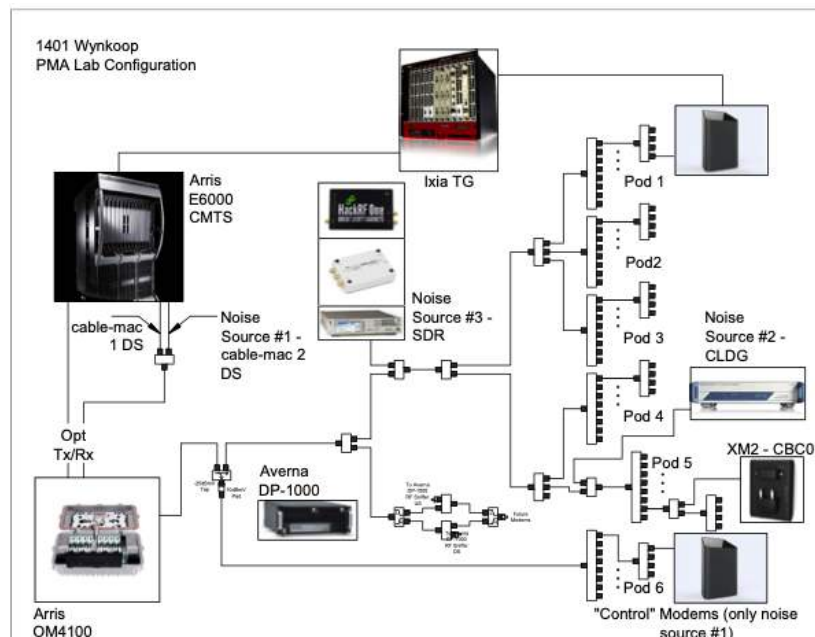
- Osmocom\_fft is a simple spectrum analyzer and capture tool that can be used to write sampled data directly to a file. Using a command such as:  
`osmocom_fft -f 1.89e9 -s 40.96e6`  
will launch the application and tune to 1.89 GHz.  
If the Osmocom utilities are not available, a GNU Radio Companion flowgraph can be created to output the data from a UHD:USRP source to a file sink block.  
The RX port of the SDR should be fitted with an antenna capable of receiving the LTE frequencies, such as a mini GSM/cellular quad-band antenna with 2 dBi of gain and 50 ohm impedance.
- Install Speed test on the UE LTE source, disable Wi-Fi, and begin a speed test. Most testing includes a downlink and an uplink speed test portion. If the energy is not seen at the expected frequency, the downlink test may not have completed or the UE may have changed bands or moved to a neighboring macro cell.
- The LTE signal should now be see on the analyzer. For live, raw IQ capture, Gqrx is a free open-source software-defined radio receiver that supports many popular radios including Airspy, rtl-sdr, HackRF, and USRP devices.





**Figure 8 - Real Time Spectrum Analyzers for IQ Capture (clockwise):  
osmocom\_fft, uhd\_fft, gqrx SDR**

- Combine the SDR TX Port 1 into the laboratory RF plant as needed to impact the desired set of modems. For example:



**Figure 9 - Combining of Noise Sources in Laboratory Environment**

- Through the splitting and combining network, different interference groups can be created. In this way the k-means based clustering algorithm that the PMA analytics engine implements can be fully tested, as multiple, custom OFDM channel profiles are supported.

## 8. Practical Application Example #2 – How SDR Has Aided Profile Management Application (PMA) Testing

Problem Statement: How can we re-create an OFDM channel impairment profile using a cable modem's reported RxMER values?

### SDR Solution Outline:

- Read in cable modem RxMER data file
- Convert the data from ¼ dB values to dB
- To create the impairment profile, set all values that are within a threshold (say 6-10 dB) of the maximum value to zero. We will be using the values outside of this “good” range to create our impairment.
- Convert the RxMER value from a ratio of power values to output voltage for the DAC
- Normalize the voltage value vector
- Use GNU Radio to stream out the data to an inverse fast Fourier transform block to convert from frequency to time domain
- Stream the data to either a file to use for later playback, or directly to the software-defined radio sink
- Inject this RF impairment into the lab environment physical plant to re-create the impairment seen at the customer premises

### Detailed Steps:

- The cable modem RxMER values can be obtained either directly from CMTS CLI or via TFTP download. The values are reported in quarter-dB, for example:

```
0x0000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x0020 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x0040 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x0060 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
.
.
.
0x0400 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x0420 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x0440 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x0460 00000000 0000B6B7 B9B6BCB6 B5BABAB4 B6B3B4B5 BAB3B8B3 BAB6B8B5 B5BBBBB6
0x0480 B6BAB5B9 B9B7BBB7 BCB4B4BA B4B6B7B9 B6B8B3BD B7B7B8B5 B8B8B8B6 B7B8B7B9
0x04A0 BAB3B8B5 B6B7B8B8 BCB7B8B8 B8BAB4B8 B7B6B9B3 B6B9B9B8 B7B9B9B9 B4BCBAB5
0x04C0 B6B8B4B7 B7BAB8B8 B3B8B7B4 B8B6B7B9 B9B4B8B5 B9BAB6B9 B4B6B4B9 B4B6B9BA
0x04E0 B6B9B3B7 BAB2B5B6 B5B5B4B6 B4B7B7BA B9B7B5BA B9B6B9B4 B6B7B7B7 BDBAB6B3
0x0500 B6B6B4B7 B6BABBB8 B3B8B5B7 B7B7B5B8 B6B6B6B8 B7B7B6BA B4B6B8B9 B7B5B8B7
0x0520 B9B6B5B8 B3B9B7B6 B6BBB7B5 B8B9B4B7 B8B7B4B7 B7B7B8BA B8B7B5B4 B4B4B9B3
.
.
.
0x0F40 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x0F60 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
```

```

0x0F80 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x0FA0 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x0FC0 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0x0FE0 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000

```

If using the values from the TFTP downloaded RxMER file, the RxMER per subcarrier data follows after a header using the specified format:

**Table 2 – RxMER TFTP File Format (4)**

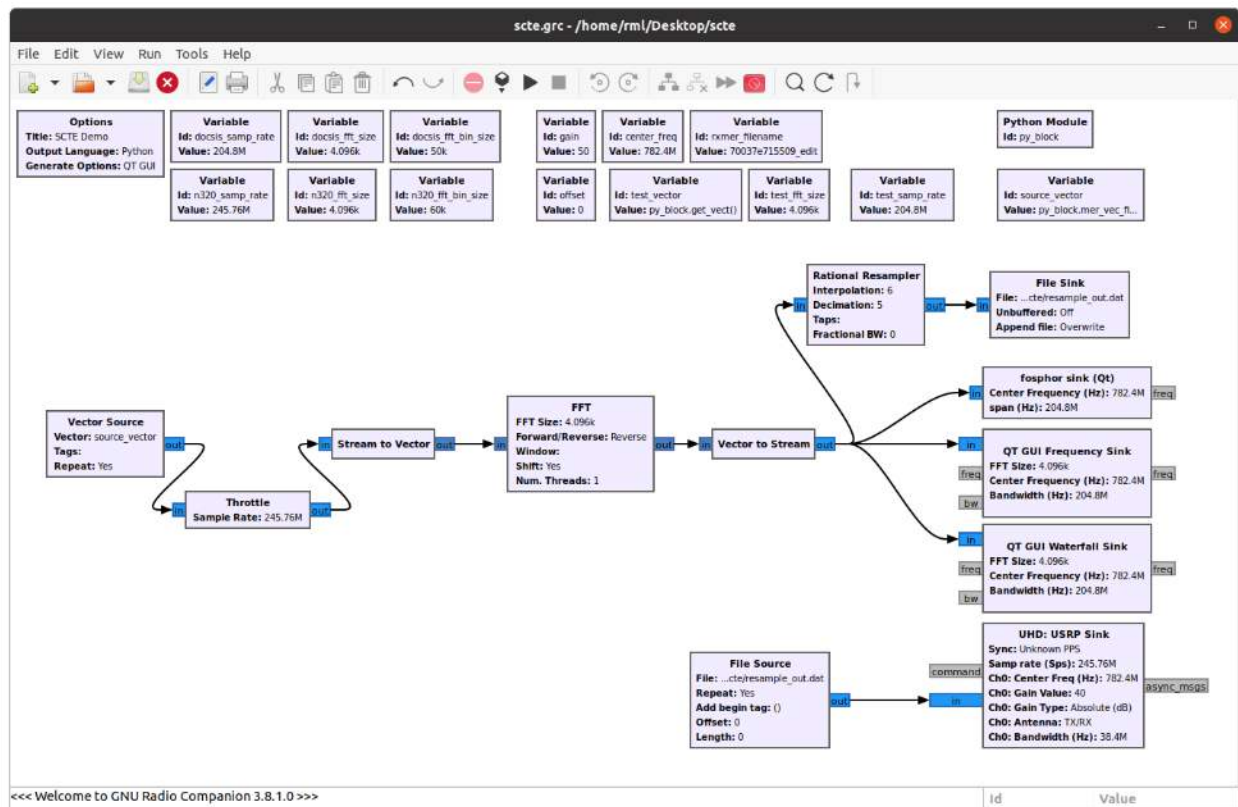
Element	Size
File type (value = 504E4E04)	4 bytes
Major Version (Value = 1)	1 byte
Minor Version (Value = 0)	1 byte
Capture Time	4 bytes
DS Channel Id	1 byte
CM MAC Address	6 bytes
Subcarrier zero frequency in Hz	4 bytes
FirstActiveSubcarrierIndex	2 bytes
Subcarrier spacing in kHz	1 byte
Length in bytes of RxMER data	4 bytes
Subcarrier RxMER data	RxMerData

The other data such as the subcarrier zero frequency, first active subcarrier index, subcarrier spacing, and length in bytes of the RxMER data will all be used to fully describe the OFDM channel impairment profile.

2. The RxMER data itself is defined as the ratio of the average power of the ideal QAM constellation to the average error-vector power. The error vector is the difference between the equalized received pilot or preamble value and the known correct pilot value or preamble value. The reported values are represented in hexadecimal values of ¼ dB. Multiple each value in the array by ¼ to properly scale.
3. In this step we discard (set to zero) all values that are within a configurable threshold of the maximum reported RxMER value. Only the low RxMER values will be used to construct the impairment vector. This example used a value of 10 dB.
4. Convert every value in the impairment array to an output voltage value for the DAC using the formula:  

$$V = 10^{(dB/20)}$$
5. Divide every voltage value by the sum of values (new values will all sum to 1) to avoid overloading the SDR output stage.
6. The processing of the RxMER data, along with the streaming to either file or directly to the software-define radio sink was achieved through the creation of a GNU Radio Companion flowgraph (Figure 10).






**Figure 10 - GRC Flowgraph to Construct and Transmit Custom Impairment Waveform**

### Flowgraph Details:

#### Embedded Python Module-

The previous six steps were all done in the embedded Python Module block. GNU Radio companion allows you to create and use custom Python code within your flowgraph. In this example, all the RxMER data manipulation and plot generation was done within the embedded Python block. After reading in the modem's RxMER data, the values are written to disk as a Matplotlib Pyplot (Figure 11):



**Figure 11 - Graph of RxMER Data Generated by Embedded Python Block**

#### Vector Source Block -

This flowgraph begins with the vector source block. This block streams items based on the input vector, which in this case is the output of a Python method in our embedded Python block.

#### Throttle Block -

This stream feeds a throttle block. This block will be set to “bypass” when using the real SDR hardware which will set its own sampling rate. It is used in combination with the software sinks or file output, when the sampling rate would only be determined by the host’s CPU clock.

#### Stream to Vector Block –

This block converts the stream of items into a stream of vectors containing N number of items. In this case the number of items will be the FFT size (4096). The stream-to-vector and vector-to-stream blocks are special in that the input and output types differ by a decimation or interpolation factor. In our example, the stream-to-vector block produces one output vector for every 4096 input samples.


#### FFT Block –

The complex valued output vector from the previous stage feeds the fast Fourier transform block. The key parameters of the FFT block are: FFT size (number of samples used in each FFT calculation), direction (forward for FFT, reverse for inverse FFT), window, shift (puts DC – 0 Hz – in the center of the FFT block for complex input type), and number of threads.


#### Vector to Stream Block –

The output of the FFT block will always be complex-valued. This block will convert a stream of vectors from the FFT block to a stream of items that can be fed directly to an output file, the

SDR hardware, a GUI frequency sink that offers spectrum analyzer functionality, or a Fospor or Waterfall sink for more real-time spectrum analyzer-like output.



**Figure 12 - Qt GUI Frequency Sink Visualization of Impairment Signal**

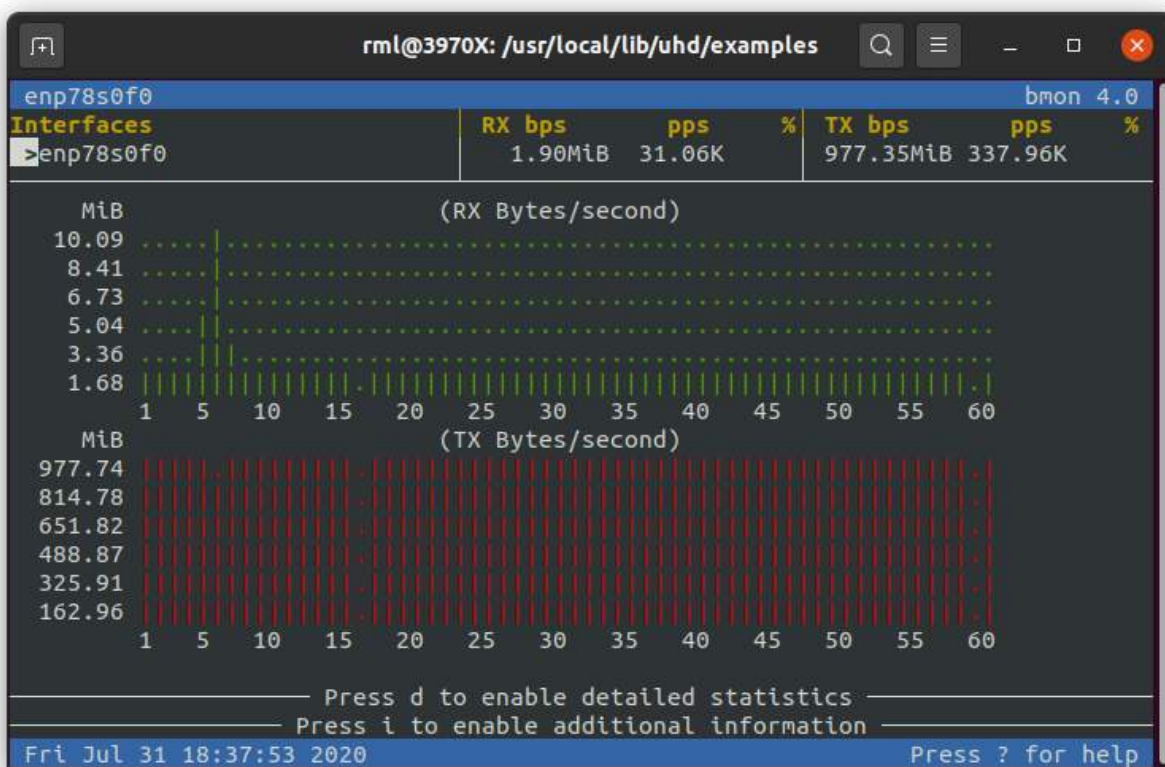


**Figure 13 - Qt Fospor Sink Visualization of Impairment Signal**

### Rational Resampler Block –


If trying to stream a file that has been sampled at a rate that is not supported by the SDR hardware, the rational resampler block allows us to change the sampling rate (e.g., from the DOCSIS standard 204.8 MHz). In this example the default hardware sampling rate is 245.76 MHz so we could set the interpolation factor to 6 and the decimation factor to 5 to upsample to 245.76 MHz.

7. If the output is sent to a file, we can replay directly to the SDR sink without continuously performing the IFFT calculation repeatedly. The act of streaming 32 bit IQ data at 245.76 MHz is very I/O intensive and requires a powerful host computer with a 10 gigabit interface. As seen in Figure 14 this stream from the host to the SDR hardware averages around ~7.8 Gbps:



**Figure 14 - Bandwith Monitor of IQ Data Stream**

8. The SDR TX output can then be combined into the laboratory plant, and the effect can then be observed in the OFDM channel RxMER data (Figure 15).



**Figure 15 - Impairment Signal on Spectrum Analyzer**

In this case we were testing the response of our profile management application (a.k.a. Octave) to verify that proper bitloading values were being set for the impacted frequency ranges. For an example OFDM profile, see Figure 16.

DS S/C/CH	DS Prof	Low-High Freq Edge (MHz.KHz)	Bits/Symbol/ Mod	Subcarrier	MaxRate Mb/sec
9/0/48	0	-default-	256qam	8.00	584
9/0/48	1	-default-	4096qam	11.73	856
9/0/48	1	702.000-708.000	2048qam	-	-
9/0/48	1	735.000-737.000	2048qam	-	-
9/0/48	1	738.000-744.000	2048qam	-	-
9/0/48	1	745.000-756.000	2048qam	-	-
9/0/48	2	-default-	4096qam	10.75	785
9/0/48	2	795.000-801.000	1024qam	-	-
9/0/48	2	802.000-813.000	512qam	-	-
9/0/48	2	816.000-828.000	512qam	-	-
9/0/48	2	831.000-832.000	2048qam	-	-
⋮					
⋮					
⋮					

**Figure 16 - OFDM Profile for SDR Generated Impairment Signal**

## 9. Equipment Used

### Software-Defined Radio Hardware:

1. Ettus Research USRP N320
2. Ettus Research USRP B210 Mini
3. Great Scott Gadgets HackRF One
4. 75 MHz to 1 GHz telescoping antenna
5. SMA Make Right Angle antenna (850/900/1800/1900/2100 MHz, 2 dBi gain, 50 ohm)
6. 10 Gb direct-attach copper cables

### Software-Define Radio Software:

1. GNU Radio v3.8.1.0
2. Ettus Research UHD drivers UHD\_3.15.0.HEAD-0-gaea0e2de
3. GNU C++ v9.3.0
4. Python 3.8.2
5. Various libraries including libosmodsr-dev, libhackrf-dev, gr-osmosdr, gr-fosphor,

### SDR Controller Host Computer:

- CPU - AMD Ryzen Threadripper 3970X
- Motherboard - Asus ROG Zenith II Extreme Alpha
- RAM – 32 GB Corsair DDR4 3600
- Storage – Samsung 970 EVO M.2 1 TB (x2)
- GPU - EVGA GeForce RTX 2080 Ti Xc Hybrid 11 Gb DDR6
- Network – Asus XG-C100C 10G Network Adapter PCI-E x4
- Power – Thermaltake 1200 W 80+ Platinum

### OS Software:

1. Ubuntu 20.04 LTS
2. Windows 10

### Spectrum Analyzer:

- Siglent SSA 3021X Plus (9 kHz to 2.1 GHz)

## 10. Conclusion

In this paper, we have explored what software-defined radio is and what it offers to engineers, students, and designers. We presented a general architecture as well as some needed background information on sampling theory and the discrete Fourier transform. Finally, the hardware, software, and drivers used to implement real data flows to sample and/or generate signals were detailed, along with two laboratory use cases.

With the abundance of open-source community supported software and low cost hardware, software-defined radio is proving to be a valuable laboratory and field testing tool. By abstracting the functionality from a specific piece of physical hardware, we can transmit and receive various modulation methods, perform signal processing and filtering, and customize routines through Python extensions all within a common software framework. These advantages, combined with the ability to embed Python code for customized signal processing routines, positions software-defined radio as a valuable laboratory and field testing tool.

# Abbreviations

ADC	analog-to-digital converter
AWS	advanced wireless service
bps	bits per second
CLI	command line interface
CMTS	cable modem termination system
CPU	central processing unit
DAC	digital-to-analog converter
dB	decibel
dBi	decibel isotropic
DC	direct current
DFT	discrete Fourier transform
DSP	digital signal processing
DTFT	discrete time Fourier transform
FDD	frequency division duplex (or duplexing)
FFT	fast Fourier transform
FPGA	field programmable gate array
GB	gigabyte
GbE	gigabit Ethernet
Gbps	gigabits per second
GHz	gigahertz
GPSDO	Global Positioning System disciplined oscillator
GRC	GNU Radio Companion
GSM	global system for mobile communications (originally <i>groupe spécial mobile</i> )
GUI	graphical user interface
Hz	hertz
IEEE	Institute of Electrical and Electronics Engineers
IF	intermediate frequency
IFFT	inverse fast Fourier transform
I/O	input/output
IQ	in-phase/quadrature
ISBE	International Society of Broadband Experts
kHz	kilohertz
LO	local oscillator
LTE	long term evolution
MHz	megahertz
MS/s	megasamples per second
OFDM	orthogonal frequency division multiplexing
OS	operating system
PCS	personal communications service
PMA	profile management application
PPS	pulse per second



QAM	quadrature amplitude modulation
QSFP+	quad small form-factor pluggable plus
RAM	random access memory
RF	radio frequency
RU	rack unit
RX	receive (or receiver)
RxMER	receive modulation error ratio
SCTE	Society of Cable Telecommunications Engineers
SDR	software-defined radio
SFP+	enhanced small form-factor pluggable
SMA	subminiature version A [connector]
SoC	system on a chip
TFTP	Trivial File Transport Protocol
TX	transmit (or transmitter)
UE	user equipment
UHD	USRP hardware driver
USB	universal serial bus
USRP	universal software radio peripheral

## Bibliography & References

1. SDRF-06-R-0011-V1.0.0: SDRF Cognitive Radio Definitions
2. <https://www.testandmeasurementtips.com/temporal-spatial-aliasing-signal-processing>
3. ETSI TS 136 508 V15.5.0 LTE; Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Packet Core (EPC); Common test environments for User Equipment (UE) conformance testing
4. Cable Modem Operations Support System Interface Specification CM-SP-CM-OSSIv3.1-I17-200610

# **Enabling Industry 4.0 Business Models For MSOs Using Wireless Mesh Networks At 60 GHz**

A Technical Paper prepared for SCTE•ISBE by

**Elliott Hoole**

Director, Wireless R&D

Charter Communications

6360 S. Fiddlers Green Circle, Greenwood Village, CO 80111

(720) 536-9424

Elliott.Hoole@charter.com

# Table of Contents

Title	Page Number
1. Introduction.....	4
2. Industrial Automation System Requirements .....	4
3. 60 GHz Indoor Deployment Aspects .....	8
3.1. Indoor Scattering .....	8
3.2. Antennas .....	8
3.3. Diversity and Coverage Robustness .....	10
3.4. Fast Reassociations .....	10
3.5. Edge Computing System .....	10
4. Performance Evaluations .....	11
4.1. First Indoor Trial .....	11
4.2. Second Indoor Trial .....	17
5. Use Cases and Demonstrations .....	18
5.1. Demonstration System #1 .....	18
5.1.1. Use Cases Demonstrated .....	18
5.1.2. Demonstration System Architecture .....	19
5.1.3. Demonstration System Data Messaging Architecture .....	20
5.2. Demonstration System #2 .....	20
5.2.1. Use Cases Demonstrated .....	21
5.2.2. Demonstration System Architecture .....	21
6. Economics and Ecosystem .....	22
6.1. Example Network Comparison.....	22
6.1.1. 60 GHz Example Network Costs .....	22
6.1.2. 5G Example System Costs .....	23
6.2. Ecosystem Comparisons .....	24
6.2.1. 60 GHz Equipment Ecosystem .....	24
6.2.2. 3GPP 5G Equipment Ecosystem.....	24
7. Conclusion.....	25
Abbreviations .....	26
Bibliography & References.....	26

## List of Figures

Title	Page Number
Figure 1 – The relationship between service availability and reliability .....	7
Figure 2 - Industrial Automation network delays.....	8
Figure 3 - Typical 60 GHz MIMO antenna array module .....	9
Figure 4 - 60 GHz indoor test network .....	11
Figure 5 - Example test network configuration.....	12
Figure 6 - Round trip delay statistics for the 1-hop test network configuration .....	13
Figure 7 - Round trip delay statistics for the 3-hop test network configuration.....	14
Figure 8 - Round trip delay statistics for the 5-hop test network configuration.....	15
Figure 9 - Measured round trip times for various packet sizes .....	16
Figure 10 - Measured round trip time with 20 and 50 packets per second.....	16
Figure 11 - Demo #1 system architecture .....	19

Figure 12 - Demo #1 data messaging architecture.....	20
Figure 13 - Demo #2 system architecture.....	21
Figure 14 - Example industrial space.....	22

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Industrial automation application areas and associated use cases .....	5
Table 2 - - Selected 3GPP technical requirements for industrial automation use cases.....	6
Table 3 - Motion control use case sub-categories .....	6
Table 4 - MIMO antenna module parameters .....	9
Table 5 - Hemi-spherical antenna module parameters.....	10
Table 6 – Measured RSSI (dBm).....	17
Table 7 - Measured round trip delay (microseconds) .....	18
Table 8 - Demo #1 use cases .....	19
Table 9 - Demo #2 use cases .....	21
Table 10 - Example 60 GHz network CapEx costs.....	23
Table 11 - Example network 5G CapEx costs .....	23

## 1. Introduction

The term “Industry 4.0” refers to the 4<sup>th</sup> industrial revolution brought about by digital technologies to dramatically increase operational efficiencies primarily through the use of analytics and automation. Wireless systems are seen as playing a primary role in this revolution to more easily allow for on-demand manufacturing process reconfigurations along with high volume data collection via industrial internet of things (IIoT) networks. The motivation of the project described in this paper is to investigate the performance and feasibility of using multi-hop pseudo-mesh indoor networks in the 60 GHz V-Band for industrial automation applications to support Industry 4.0 use cases. These use cases include the real-time control of multiple robotic apparatuses using a dedicated multi-access edge compute (MEC) system. Another goal of the project is to characterize the deployed wireless connection robustness and their practical limitations in both static and dynamic links. In addition to performance evaluation it is also important to understand the current state of supporting component technologies and the viability of a robust technology ecosystem for deploying and supporting these networks.

The main challenge for the initial phases of this project was the fact that there are no commercially available 60 GHz true mesh products currently available on the market. Mesh connectivity for this project was emulated using multiple point-to-multipoint links. Another challenging aspect was that the 60 GHz Consumer Premises Equipment (CPE) endpoints in our trials are the same physical devices as the access nodes with a different operational configuration which would not be the case in an actual commercial deployment. Lastly there were some use cases on our trials which required connectivity translation to interwork with 60 GHz network, so there may be issues addressing some of the needs of some potential customers depending on their specific objectives and equipment. Many of these issues are being addressed in later phases of this project and will be detailed in a future publication.

However despite the challenges of these initial project phases, it has been found through measurements and practice that networks with 60 GHz technology can meet needs of (Industrial Automation) (IA) use cases today if deployed and managed properly.

## 2. Industrial Automation System Requirements

There are many specific situations within the Industrial Automation use case umbrella. One expert who is active in the field has stated that a good place to start would be to “duct-tape an iPhone to a milling machine”[1] which would enable the collection of a production asset’s operational parameters. A very high value use case that does not require huge bandwidth and low latency is to perform what is known as “finding the hidden factory”. This refers to operational inefficiencies which are found through process monitoring, data collection, and targeted analytics to determine the efficiency of the end-to-end production process and identify sections of it that may not be performing as intended. In so doing, surprisingly large amounts of money are saved simply by reducing otherwise unknown waste [2].

The 3rd Generation Partnership Project (3GPP) have identified target requirements for specific IA use cases, and these can be found in TS 22.104 [3]. **Error! Reference source not found.** below is a table showing the 3GPP use cases (columns) with their associated application areas (rows). As seen, not every application area employs every use case. So different industrial customer segments can potentially be addressed with different target service and application packages, although this concept is not explored further in this paper.

**Table 1 - Industrial automation application areas and associated use cases**

3GPP TR 22.104 Annex A	Motion Control	Control-to-Control	Mobile Control Panels With Safety	Mobile Robots	Remote Access and Maintenance	Augmented Reality	Closed Loop Process Control	Process Monitoring	Plant Asset Management
Factory Automation	X	X		X					
Process Automation				X			X	X	X
HMI's and Production IT			X			X			
Logistics and Warehousing		X		X					X
Monitoring and Maintenance					X				

Each of the use cases listed in Table 1 has a set of system level parameters which should be met in order to provide the expected level of performance. A representative sampling of these requirements is given below in Table 2.

**Table 2 - - Selected 3GPP technical requirements for industrial automation use cases**

Use Case	Service Avail	Message Size	Transfer Interval	Survival Time	UE Speed	# of UEs	Service Area
Motion Control	5-7 9's	50 bytes	0.5 ms	0.5 ms	≤ 75 kph	≤ 20	50x10x10 m
Mobile Robots	6 9's	40 - 250 bytes	1-50 ms	1-50 ms	≤ 50 kph	≤ 100	< 1 sq km
Process Monitoring	4 9's	20 – 255 bytes	100 - 60k ms	3 x Trans Intv	0	10k – 100k	10x10x0.05 km
Mobile Control Panel	6-8 9's	40 - 250 bytes	4-8 ms	4-8 ms	≤ 8 kph	TBD	50x10x4 m
Process Control	6-8 9's	20	10 ms	0	0	10-20	100x100x50 m
Control-to-Control	6-8 9's	1k bytes	10 ms	10 ms	0	5-10	100x30x10 m
Augmented Reality	3 9's	unspec	< 10 ms	unspec	< 8 kph	unspec	20x20x4 m
Asset Management	4 9's	20 – 255 bytes	'several seconds'	3 x Trans Intv	0	Up to 100k	10x10x0.05 km

In addition to the requirements listed in Table 2, 3GPP TR 22.104 [3] further segments each use case into sub-categories so that there are multiple sets of requirements for each of the use cases listed in the columns of Table 1. For example, the Motion Control use case (UC) has 3 sub-categories which are listed below in Table 3.

**Table 3 - Motion control use case sub-categories**

Use Case	Service Avail	Message Size	Transfer Interval	Survival Time	UE Speed	# of UEs	Service Area
Motion Control UC #1	5-7 9's	50 bytes	0.5 ms	0.5 ms	≤ 72 kph	≤ 20	50x10x10 m
Motion Control UC #2	6-8 9's	40 bytes	1 ms	1 ms	≤ 72 kph	≤ 50	50x10x10 m
Motion Control UC #3	6-8 9's	20 bytes	2 ms	2 ms	≤ 72 kph	≤ 50	50x10x10 m

The most stringent use case sub-category of all those listed in [3] is Motion Control #1. This is the use case that is typically focused on in marketing literature and used as *the* requirements for Industrial Automation, but as can be seen in Table 2 there are other sub-categories of Motion Control that are not as

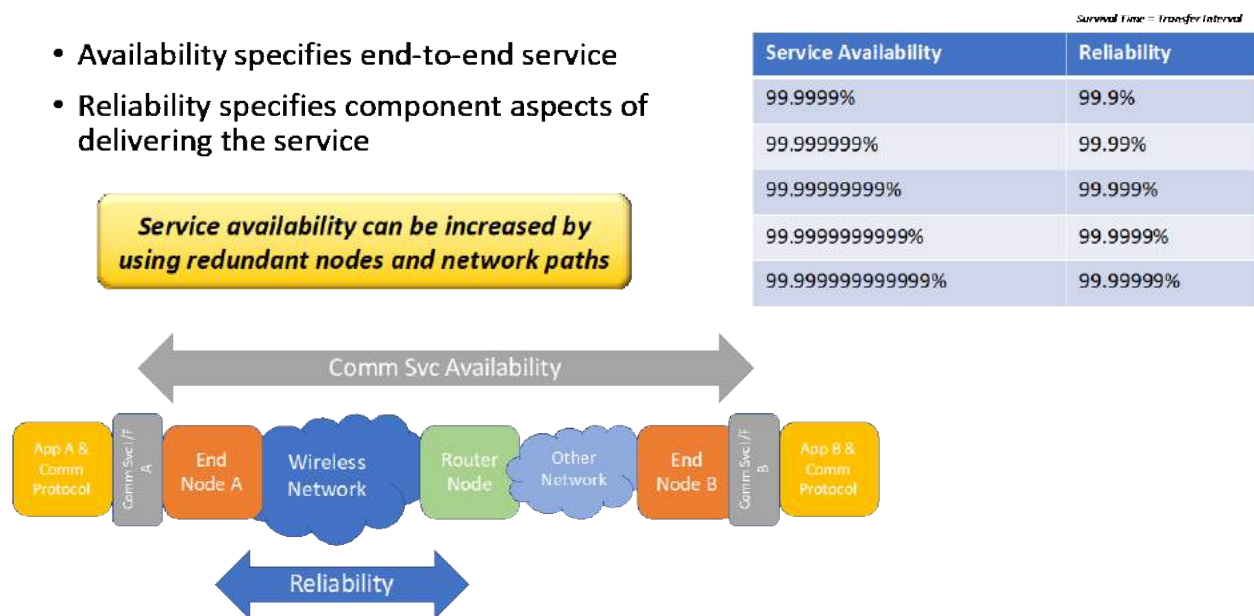
demanding as Case #1 and many other use cases and their sub-categories which are not as demanding as Motion Control #1 but are very useful for Industrial Automation networks and services.

Consider a 100 byte packet sent every 1 msec. This results in a required throughput of  $8 \times 100 \times 1000 = 800$  kbits/sec. As can be seen from the previous figures this represents a sufficient amount of bandwidth and transmission interval for every listed use case above. So a system throughput budget of 1 Mbps per Industrial Automation endpoint can be viewed as being representative of what IA systems require.

Another important requirement for IA networks is service availability which is an end-to-end network requirement for delivering specified services. System redundancies, both at the node level and in available paths, can be employed to increase the overall end-to-end availability beyond what a single node/path can deliver.

An illustration of service availability with respect to redundancy is shown below in Figure 1. In this case reliability concerns only the wireless network portion of the end-to-end network. The table embedded in Figure 1 gives the relationship between service availability and reliability when the survival time is equal to the transfer interval. These will be discussed shortly. More information concerning service availability and its relationship with reliability can be found in 3GPP TS 22.241[4].

In order to increase the reliability of the wireless network portion of Figure 1 and the overall availability of the end-to-end service, a wireless mesh network can be employed. This type of network provides an endpoint with multiple paths through which to send data so as to eliminate the sole dependency on any particular node for communication.

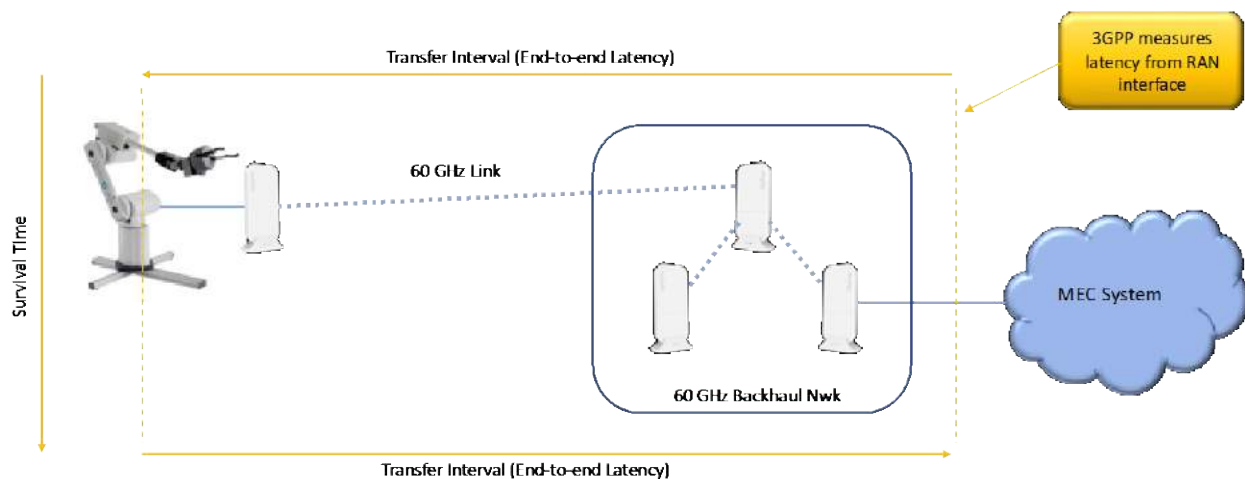


**Figure 1 – The relationship between service availability and reliability**

As seen in Figure 1, service availability and reliability are related by parameters called transfer interval and survival time. These are system delays that affect an endpoint's response time to a command from a control system. These delays are illustrated below in Figure 2. In this illustration the transfer interval from the MEC to the endpoint is the same as the one from the endpoint to the MEC since they traverse the



same path. This is not always the case and must be taken into account for round trip response time calculations. The survival time is the amount of time it takes an endpoint to respond to a command and issue a status response back to the control system.



**Figure 2 - Industrial Automation network delays**

The IA network requirements discussed in this section illustrate the level of performance expected to be able to perform the target IA use cases. As seen in Table 2 not every use case requires a very high degree of performance and many of them can be enabled with equipment that is widely available.

### 3. 60 GHz Indoor Deployment Aspects


#### 3.1. Indoor Scattering

Historically 60 GHz wireless transmission has been regarded as being problematic due to the oxygen absorption characteristic in the radio frequency band. This phenomenon really only affects very long links as the effect is described in dB/km. However for shorter links (< 100m) this effect is negligible.

60 GHz radio signals do not pass through materials as well as radio signals do in other bands. As such there is a tendency for a large amount of environmental scattering which makes antenna directionality indoors much less critical or even important. For indoor 60 GHz applications this means that line-of-sight and precise beam pointing deployment aspects are not as critical as in other bands where the radio signals scatter less.

#### 3.2. Antennas

The typical antenna devices in 60 GHz products is found to be a patch array that is used in conjunction with beamforming/beam-switching to provide higher gain for longer links and alleviate the need for antenna alignment during installation. The standard antenna array module is similar to the one shown below in Figure 3.



**Figure 3 - Typical 60 GHz MIMO antenna array module**

These standard multiple-input and multiple-output (MIMO) patch array modules have the typical specifications given below in Table 4.

**Table 4 - MIMO antenna module parameters**

MIMO Module Parameter	Value
Number of elements	64
Coverage angle	90° horizontal / 40° vertical
Max Effective Isotropic Radiated Power (EIRP)	36 dBm

In a representative warehouse indoor environment we have seen through field testing, which will be further described in Section 4.2, that a hemi-spherical pattern from the same AP is better overall in terms of coverage and robustness. This was done by substituting a hemispherical antenna module for the standard MIMO array module and then testing both configurations in the same deployment locations. The hemispherical antenna has the specifications given below in Table 5.

**Table 5 - Hemi-spherical antenna module parameters**

Hemi-Spherical Module Parameter	Value
Number of elements	32
Coverage angle	180° horizontal / 180° vertical
Max EIRP	30 dBm

For product commercialization, the radio units with hemi-spherical antenna modules can simply be variants of existing products with an antenna module substitution.

### **3.3. Diversity and Coverage Robustness**

In order to meet the more stringent requirements for IA use cases (e.g. motion control), deployments should make use of overlapping coverages from more than one access point (AP). During the network planning phase, the coverage areas from different APs should be planned with substantial overlap to allow for secondary or alternate connectivity options for each endpoint.

One possible system configuration could be to install multiple access point radio units feeding a single baseband unit which would then employ various diversity combining techniques. This concept is being investigated in an ongoing phase of this project.

### **3.4. Fast Reassociations**

IEEE specification 802.11r-2008 [5] describes a fast transition technique for a client station (STA) to move its radio connection from one AP to another in a more expedited manner without having to undergo the full authentication procedure [7]. Doing so can reduce the reassociation time to perhaps 50 msec and is recommended to be done within enterprise networks using WPA2 Enterprise security. This technique should be utilized in IA systems to minimize the impact of reassociations on system performance.

It may be possible to make use of an edge compute system to help facilitate these transitions. This concept is being investigated in an ongoing phase of this project.

### **3.5. Edge Computing System**

To truly enable the most meaningful IA use cases requires an edge compute system. A computing system is required to provide control and computing functionality for automated industrial endpoints and tasks. Having this computing system collocated with the endpoints minimizes latency and enables certain use cases that could not be enabled if the system was located further back in the network.

## 4. Performance Evaluations

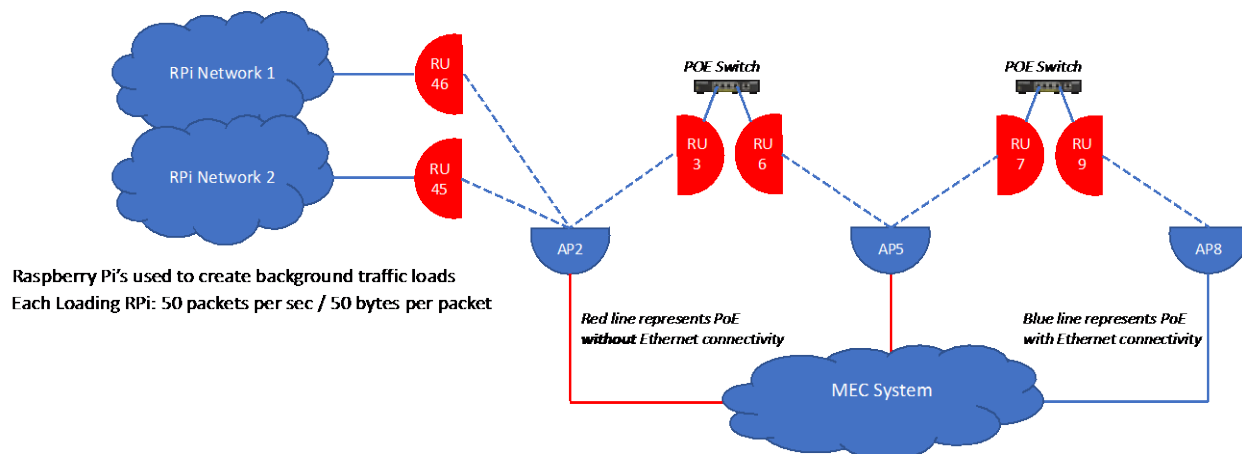
### 4.1. First Indoor Trial

In order to test and evaluate the suitability of 60 GHz for use in Industrial Automation networks, we constructed an indoor network consisting of 10 nodes that were installed around the floor of a typical office building environment. Another 4 units were nomadic meaning their location was not permanent and could range throughout the coverage area. All of the radio nodes used in this trial were from the same vendor, identical, and based on a well known 802.11ad chipset. Three of the units were configured to be point-to-multipoint access points, and the rest were configured as client devices of those access points. Along with the radio nodes, a server cluster was installed in the telco room on the same floor to serve as a MEC system. The layout of the installed network is shown below in Figure 4.



**Figure 4 - 60 GHz indoor test network**

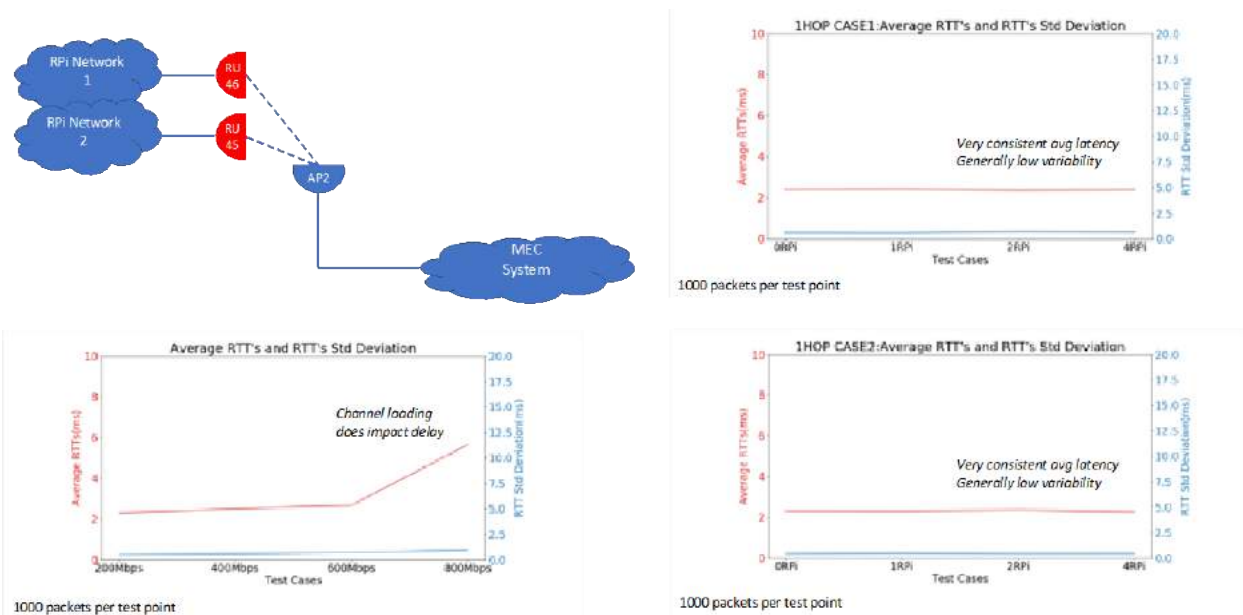
Since the deployed radio units do not provide a true mesh operation capability, we needed to emulate one by using multiple point-to-multipoint connections which could be configured and controlled through the switching infrastructure. To provide characteristic industrial network loading we used multiple Raspberry Pi devices to emulate industrial endpoints with 50 byte payloads at 50 packets per second. A representative test network configuration is shown below in Figure 5.



**Figure 5 - Example test network configuration**

In the example test network of Figure 5 the blue half-circles represent APs and the red half-circles represent remote units (RUs). The only egress point for packets from the MEC system is the blue line to AP8. The red lines to AP5 and AP2 have PoE to power the access points but the Ethernet connection on each switch port has been disabled. So the packets from the MEC system must traverse the path from AP8 to RU46 through 5 airlink hops. This effectively emulates a situation where the traffic from an endpoint would traverse multiple intermediate mesh nodes before arriving at the intended destination.

Several network configurations were employed and round trip delay measurements were collected for different scenarios with different types of traffic loading. In the cases labeled “Case1”, the device under test (DUT) was connected to the same RU as the Raspberry Pis that were emulating industrial endpoints and a varying number of Raspberry Pis were used. In the cases labeled “Case2”, the DUT was connected to a different RU than the Raspberry Pis but both RUs were connected to the same AP. Lastly bidirectional iperf traffic was used in differing amounts to create network loading. The results are shown in Figures 6-8.

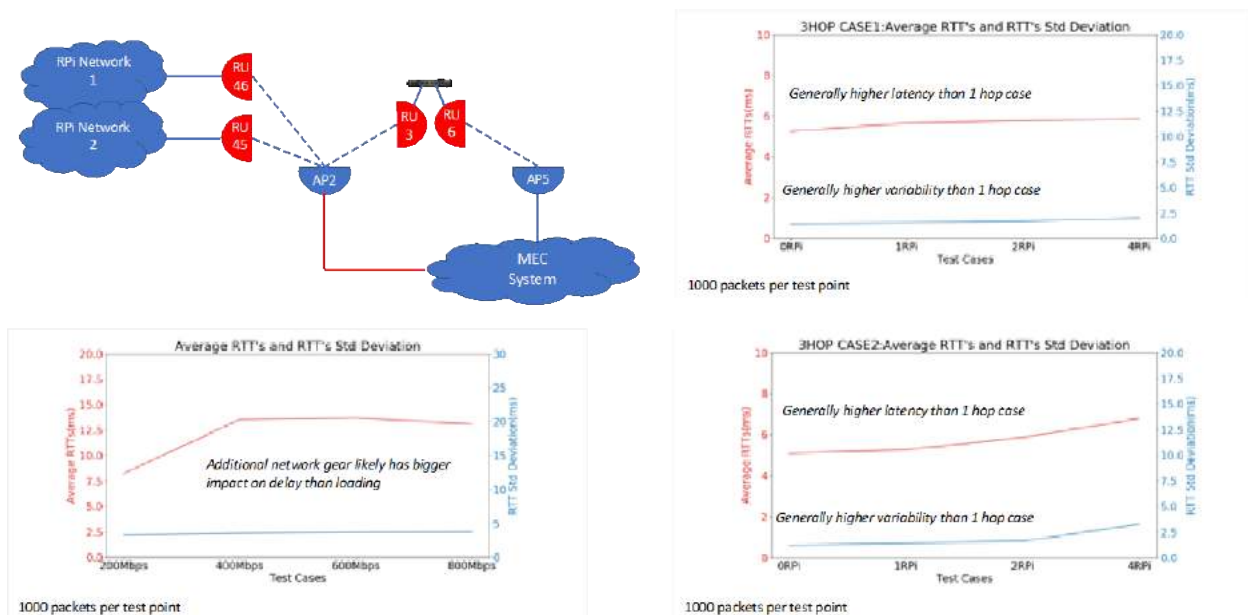


**Figure 6 - Round trip delay statistics for the 1-hop test network configuration**

As seen in Figure 6 the two sets of data from the Raspberry Pi cases are very similar. There does not seem to be enough loading from the endpoint emulators to really affect the outcomes significantly. This means that in a scenario with a fairly low number of industrial endpoints per access point with a direct connection to the MEC system the delay performance should be very consistent. This is further corroborated by the 3<sup>rd</sup> case with the iperf traffic. The delay results are very similar to the Raspberry Pi cases until the background traffic gets very high at which point the delay jumps significantly. It has been observed with certain chipsets that airlink errors which require retransmissions can stall the data queue and create a buffering delay which impacts the transmission time. This effect appears to be the cause of the increased delay for the last iperf case data point. But with lower amounts of background traffic the results track those of the Raspberry Pi cases.

In the test network configuration with 3 wireless hops as shown in Figure 7, the Raspberry Pi case results are largely just an overall increase from the 1-hop case due to the increased amount of network elements. However the last data point in Case 2 seems to show a slight increase from the buffering delay effect mentioned above. The iperf background data results show a more dramatic variation, likely due to the increased number of airlink hops presenting more opportunities for retransmissions and thus a larger impact from the buffering delays.



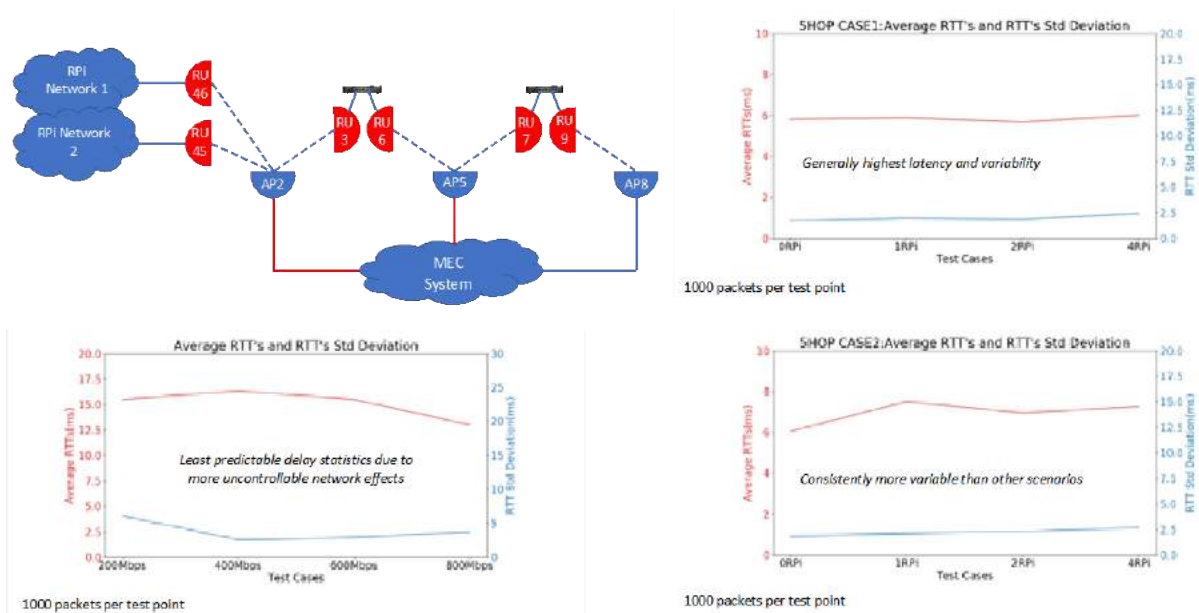


**Figure 7 - Round trip delay statistics for the 3-hop test network configuration**

For the test network configuration with 5 airlink hops, the results are even more varied as is expected. The Raspberry Pi cases are still generally consistent with more overall delay due to the increased amount of network. Again Case 2 shows more variation than Case 1 with some buffering delays being a likely culprit for the increased delay statistics during the test run.

Perhaps the most interesting results come from the iperf background data case. It is seen in that the test point with 800 Mbps background traffic, the mean delay is significantly lower than that of the test point with 400 Mbps background traffic. I believe this further exemplifies the buffering delay effect as there are now even more opportunities for retransmission events and the results reflect that.


A new aspect of 802.11ay that could be very beneficial in regulating the flow of traffic through the network is the time division duplex (TDD) scheduling feature. 802.11ad uses the typical listen-before-talk (LBT) mechanism which can lead to congestion and collisions at various points of the network, particularly the APs. With TDD scheduling each endpoint has a timeslot on the airlink reserved for its use along with an assigned priority. This new mechanism and its possible benefits will be investigated in an upcoming phase of this project.




**Figure 8 - Round trip delay statistics for the 5-hop test network configuration**

Other considerations that have been observed to affect delay statistics are the sizes of the packets and the number of packets per second required for an endpoint as seen below in Figure 9 and Figure 10. The data in both of these figures was captured with the network in the 5-hop configuration along with 500 Mbps of background traffic.





**Figure 9 - Measured round trip times for various packet sizes**




**Figure 10 - Measured round trip time with 20 and 50 packets per second**

The data runs for both figures was comprised of 1000 packets. The packet rate for Figure 9 was 10 packets per second. The packet size for Figure 10 was 50 bytes. As can be seen from the figures, packet size and packet frequency can both have significant effects on the measured round trip times, especially in the multi-hop network configurations. These parameters need to be carefully considered when deploying networks for industrial automation.

## 4.2. Second Indoor Trial

In another trial at a different location, 60 GHz indoor nodes from a different vendor were modified and deployed with different antenna modules to evaluate their suitability for IA applications. Hemispherical antennas were determined to be the best choice at both APs and clients. The use of hemispherical antenna modules are recommended for indoor IA applications which can provide performance improvements compared to the multiple input multiple output (MIMO) antenna modules.



**Table 6 – Measured RSSI (dBm)**

Row	A	B	C	D	E	F	G
0	-55	-59	-57	-57	-60	-55	-59
1	-55	-55	-61	-55	-59	-55	-58
2	-59	-57	-58	-54	-60	-55	-57

**Table 7 - Measured round trip delay (microseconds)**

Row	A	B	C	D	E	F	G
0	282	333	397	380	345	339	311
1	390	363	351	468	452	471	532
2	410	384	354	385	418	369	380

Delay measurements seen in this second trial are much lower and more consistent, so equipment selection is critical for IA network performance. Even though vendor #2 uses the same chipset as vendor #1, the performance characteristics, especially delay, are substantially different due to the specific equipment implementations. As with any network application, equipment selection should be carefully considered for IA networks to achieve the required level of performance.

Use cases that require endpoint mobility along with real-time and/or constant control cannot readily achieve the required QoS with standard 802.11ad equipment. Of course there are many factors that govern this situation, but the amount of time required for endpoint reassociations between APs can easily exceed the required control packet interval resulting in performance issues for the more demanding use cases such as motion control. Further enhancements are needed for mobility situations and robustness against connection issues due to objects moving within the environment if real-time control is required. Improvements for these scenarios are being investigated as an ongoing phase of this project.

## **5. Use Cases and Demonstrations**

In order to assess and characterize the performance of 60 GHz networks for industrial automation applications, several demonstrations were developed each employing several IA use cases.

### **5.1. Demonstration System #1**

The first demonstration involved a user-controlled mobile robot car that is piloted through a series of gates in a timed course. A live video stream from the robot car is displayed on a computer monitor and the user is told by the MEC system which gate to navigate to next. The position of a game controller joystick is read by the system and converted to servo motor parameters which are sent to the robot car to control its speed and direction. The user is then able to pilot the car using the displayed video and the joystick control. The elapsed time is also displayed for the user. Infrared (IR) sensors in each gate detect the presence of the robot car and the user is then given the next gate in the course on the screen by the system. If the wrong gate is detected by the system, this notification is displayed on the screen. The course continues until the last gate is reached at which point the elapsed time is recorded.

#### **5.1.1. Use Cases Demonstrated**

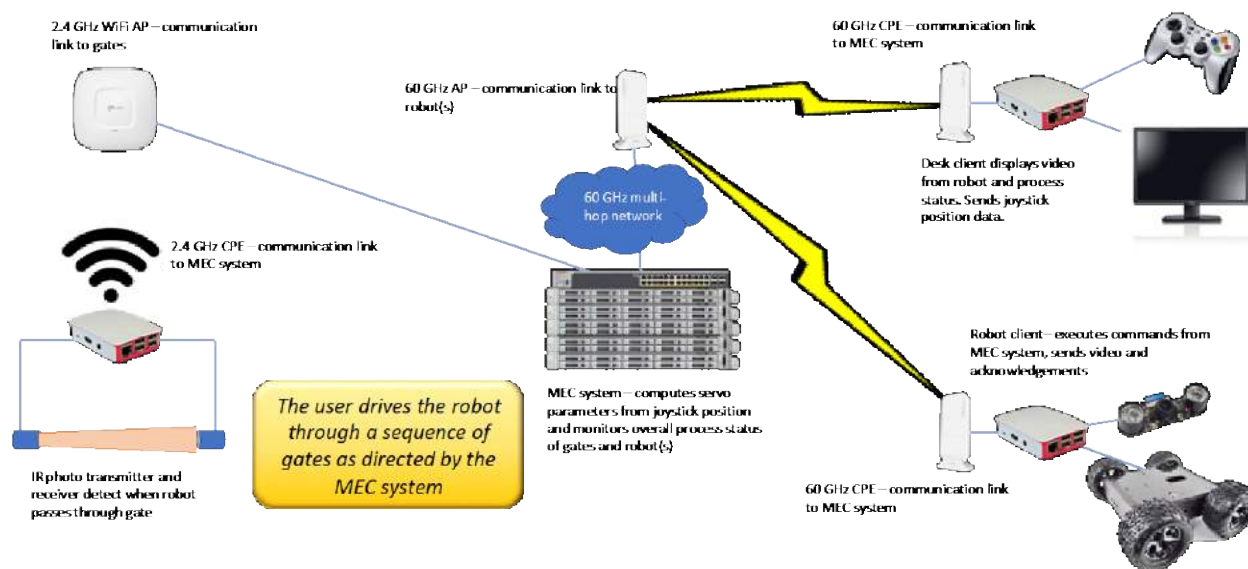
The industrial automation use cases demonstrated in Demo #1 are listed below in Table 8.

**Table 8 - Demo #1 use cases**

Use Case	Description
<b>Motion Control</b>	Robot servos updated 20 times per second according to the joystick position
<b>Control-to-Control</b>	Live video stream (4 Mbps) from robot to desk client
<b>Closed Loop Process Control</b>	Next gate indication and elapsed time display
<b>Process Monitoring</b>	Gate status messages including wrong gate detection
<b>Mobile Control Panel</b>	Real time progress and status displayed on desktop

### 5.1.2. Demonstration System Architecture

The system architecture of the demonstration system is shown below in Figure 11.



**Figure 11 - Demo #1 system architecture**

This demonstration system features a mobile robot which is controlled by a game controller joystick. The mobile robot includes a camera which streams video back to the user. The user then moves the joystick to control the speed and direction of the robot. With this interface the user directs the car through a series of gates which is monitored by the MEC system. The elapsed time is displayed on the output screen, and the overall goal is to have the robot go through the correct sequence of gates in the shortest amount of time.

### 5.1.3. Demonstration System Data Messaging Architecture

The sequence of messages that flow between the various functional entities of Demo #1 is shown below in Figure 12.

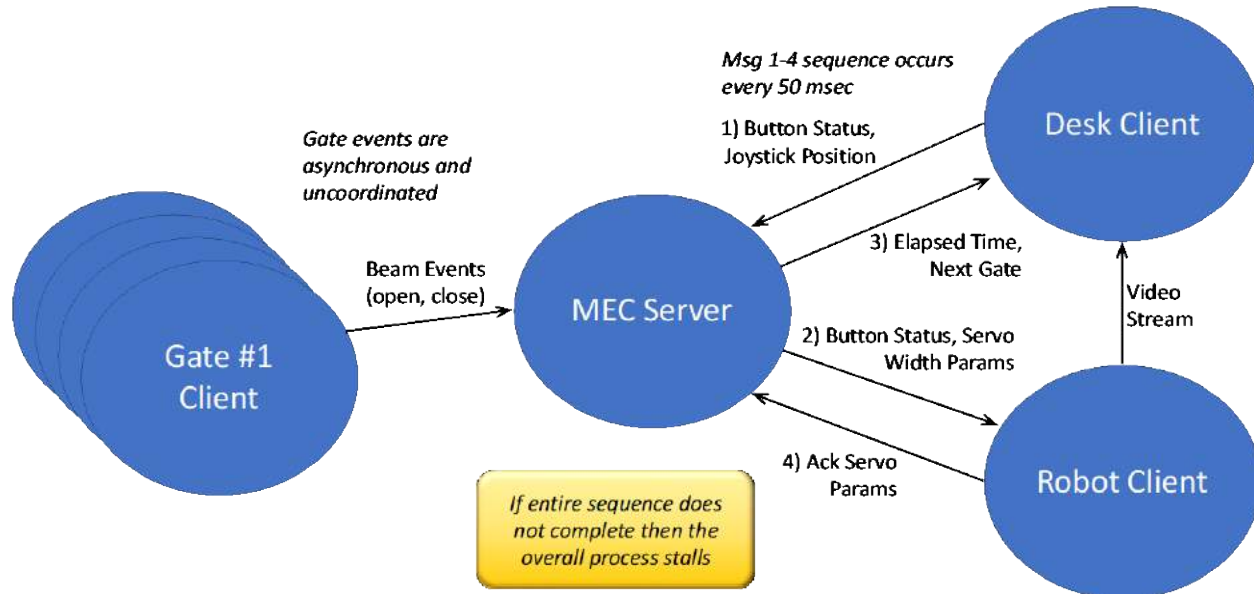


Figure 12 - Demo #1 data messaging architecture

The data architecture and message flows depicted in Figure 12 is implemented rather poorly on purpose. If any of the messages in the continuous sequence 1-4 are lost then the entire process will stall. Admittedly this is poor software design, however for our purposes this will readily illustrate when a packet loss creates a critical problem for a use case. During the normal course of the demonstration this situation never occurred. It can occur if the robot car is driven beyond the coverage area of the access point. This would be considered a mobility situation and as stated above, mobility enhancements are being investigated in an ongoing phase of this project.

## 5.2. Demonstration System #2

A second system was developed to demonstrate a fully autonomous scenario involving 2 robots performing interactive tasks as directed by the MEC system. The demo consists of a robotic arm loading a number of widgets into bins mounted on a mobile robot. The MEC system directs the mobile robot to go to the first loading position and monitors its progress during the trip. Once the mobile robot is in the correct position the MEC system directs the robot arm to load widgets into the first bin mounted on the mobile robot. The robot arm movements consist of a series of poses that are directed and monitored by the MEC system and require quite a large number of messages to perform. Once the first set of widgets is loaded, the MEC system directs the mobile robot to turn around so a second set of widgets can be loaded into a second bin. Once the mobile robot is in the new position the robot arm performs a similar series of actions as before to load another set of widgets into the other bin. Once the second set of widgets is loaded the mobile robot is sent off to its next location and its progress is monitored.

### 5.2.1. Use Cases Demonstrated

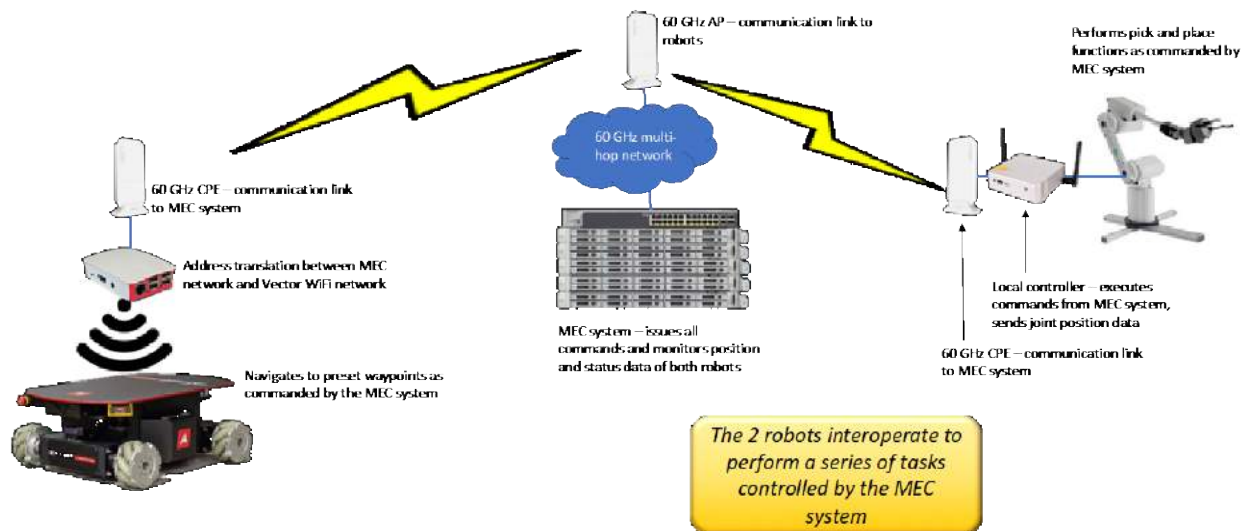
The industrial automation use cases demonstrated in Demo #2 are listed below in Table 9.

**Table 9 - Demo #2 use cases**

Use Case	Description
<b>Mobile Robots</b>	Mobile robot with 3 pre-programmed waypoints and position status checks every 100 msec. ~1100 messages exchanged during the demo.
<b>Motion Control</b>	Robot arm with 58 separate poses controlled by the MEC system. ~144,500 messages exchanged during the demo.
<b>Process Monitoring</b>	5 interlinked segments monitored and coordinated by the MEC system during the demo.

### 5.2.2. Demonstration System Architecture

The system architecture of the second demonstration system is shown below in Figure 13.



**Figure 13 - Demo #2 system architecture**

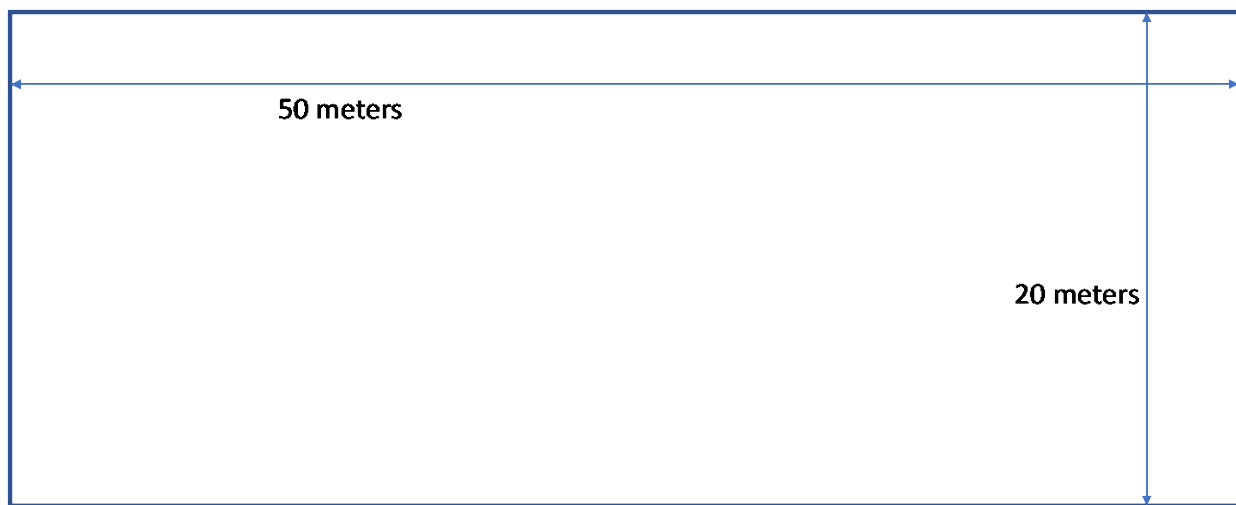
In this demonstration everything is controlled and monitored by the MEC system. The robots are programmed to perform certain tasks which requires both to operate appropriately in order to complete the overall job successfully. The overall job is for the mobile robot to collect a number of widgets which

are loaded onto it by the robot arm. So there are five distinct job phases that are controlled and monitored by the MEC system.

## 6. Economics and Ecosystem

In this section I will define and illustrate an example industrial space for comparative purposes (50 m x 20 m). The following sub-sections give the expected equipment for deploying networks based on 60 GHz nodes vs a 5G system.

### 6.1. Example Network Comparison



**Figure 14 - Example industrial space**

#### 6.1.1. 60 GHz Example Network Costs

Based on the network deployed in the 2<sup>nd</sup> indoor trial described above, it is reasonable to assume that 2 rows of 4 access points would be needed to cover the example space defined above. Although fewer APs might serve to provide adequate coverage, this number will provide the desired overlapping coverages to allow endpoints to reassociate as required to meet the necessary level of performance.

The costs of capital expenses (CapEx) for a 60 GHz network to address the example space are listed below in Table 10. These include both equipment and installation costs.

**Table 10 - Example 60 GHz network CapEx costs**

Element	Unit Cost	Quantity	Subtotal
Access Points	\$500	8	\$4000
AP Installation	\$250	8	\$2000
STA Endpoints	\$200	20	\$4000
STA Installaion	\$50	20	\$1000
Edge Compute System	\$5000	1	\$5000
Total 60 GHz (CapEx)			\$16000

The example 60 GHz network would not require any additional components beyond what is already part of the customer's IT infrastructure. Operational expenses for the network equipment (e.g. power) would be covered by the customer. The CapEx costs would be factored into a services offering by the MSO as part of a larger bundle which could include value-add services such as analytics and predictive maintenance. The relatively low network expenses, both capital and operational, could make this approach quite attractive for MSOs and their service integrator partners.

### **6.1.2. 5G Example System Costs**

For the example space the most appropriate 5G radio units would be industrial picocells. Based on the dimensions of the space it is likely that three units would be required to allow sufficient coverage overlap. Assuming the 5G core network (5GC) is located in the multiple-system operator (MSO) network, an on-premises user plane function (UPF) unit is required to handle local Ethernet traffic between the MEC and the picocells.

**Table 11 - Example network 5G CapEx costs**

Element	Unit Cost	Quantity	Subtotal
<b>Industrial Picocells</b>	\$5000	3	\$15000
<b>Picocell Installation</b>	\$250	3	\$750
<b>Enterprise UPF Unit</b>	\$2000	1	\$2000
<b>STA Endpoints</b>	\$250	20	\$5000
<b>STA Installation</b>	\$50	20	\$1000
<b>Edge Compute System</b>	\$5000	1	\$5000
<b>Total 5G (CapEx)</b>			\$28750

The 5G system does require connectivity to a 5GC which for this exercise could be located in the MSO network. The MSO would likely factor the operation and maintenance cost of the 5GC into the offered



services package which would be an additional monthly operational expense for the customer. If the customer requires or does not wish for their internal traffic to be exposed to the MSO they could be provided with an on-site 5GC to enable a private network for an additional CapEx cost plus it is likely the MSO would charge an operations and maintenance charge for the local 5GC.

Another aspect of the 5G system is that of usable spectrum. Currently 3GPP 5G systems require licensed spectrum over which to operate. The standards for 5G NR in unlicensed spectrum (NR-U) are expected to be ratified in September, 2020, and this will allow for operation of 3GPP 5G systems in unlicensed frequency bands once equipment is commercially available in 2021. However, most 5G networks are non-standalone (NSA) meaning they require a 4G anchor channel in order for the 5G channel to operate. As networks evolve these limitations will go away, but they will be a major deployment consideration for the next few years and should be factored into any service offerings and target markets.

## **6.2. Ecosystem Comparisons**

### **6.2.1. 60 GHz Equipment Ecosystem**

Currently the 60 GHz equipment ecosystem is not large. In 2018 the global millimeter wave technology market was \$289.2M with the frequency range 57-86 GHz having the primary revenue generation for the market [10]. This market is primarily focused on outdoor equipment with few offerings for indoor gear. The 802.11-based 60 GHz equipment market is expanding due to Facebook's Terragraph initiative, but that equipment is largely targeted for outdoor access and backhaul. Indoor is becoming of more interest, but it will take time for new players and offerings to come to market.

The largest 802.11ad chipset supplier is Qualcomm, but there are several others as well. Most notably Siervs and Peraso supply 802.11ad radio ICs along with Peraso and Blu Wireless who supply 802.11ad baseband ICs. Intel produces 802.11ad ICs for endpoints, but they seem to be used only in Intel modem products. With the advent of 802.11ay it remains to be seen how the chipset and corresponding equipment landscapes may change. Qualcomm is producing an 802.11ay chipset that is being used extensively in the Terragraph equipment, but it is not yet available in mass market quantities. Peraso has also announced an 802.11ay chipset and others will as well as the equipment market develops further.

Much of the currently available equipment has similar specifications due to limited number of chipset suppliers and equipment manufacturers vying for largely the same target markets. Differentiation and competitive advantages seem to come largely from support systems e.g. management, configuration, etc.

***Bottom Line: The 60 GHz chipset and equipment markets comprise both large and small players, but the equipment target markets have not taken off so overall there is not a lot of muscle behind the current offerings.***

### **6.2.2. 3GPP 5G Equipment Ecosystem**

The 3GPP ecosystem is very robust, but sadly the 3GPP ecosystem is no longer as diverse as it once was due to many consolidations amongst vendor companies. Last year the 5G global infrastructure market was nearly \$1B and by 2026 it is predicted to be over \$50B [11]. The market is supported by heavyweight equipment and chipset vendors that can be considered telecom institutions. The variety of offerings is staggering with a range of radios that run a gamut from ones that support very large cells to ones used in very small cells in a wide variety of frequency bands and combinations of frequency bands. The radio platforms can be based on anything from standard chipsets to purpose-built designs to commercial off-the-shelf (COTS) hardware.

The 5G standards targeting industrial automation operation are still expected to be fully ratified in September of 2020. This operation is called ultra-reliable low-latency communication (URLLC), and equipment that supports this feature is expected to be available near the end of 2021 at the earliest.

The 3GPP 5G standardization activities are very active and highly dynamic with 370 total 3GPP members and 1267 5G related standards in Release 16. The 5G promotional activities as well are highly spirited with many companies touting the virtues of the new technology in creative and interesting ways.

***Bottom Line: The 3GPP 5G chipset and equipment markets are dominated by heavyweight telecom giants offering a wide selection of gear for many target markets. This includes Industrial Automation with the new URLLC feature once equipment is available late next year.***

## 7. Conclusion

60 GHz radio technology can provide robust connectivity for indoor environments. Line of sight is not necessarily required due to the tendency for the radio signals to undergo a considerable amount of scattering from objects and surfaces in the environment. As such, precise antenna orientation does not seem to be as critical for the target environments. There is naturally some variability in the radio signals, but generally the radio links have been seen to be very stable and provide robust connectivity. The emerging 802.11ay standard will provide true mesh capabilities which will further enhance performance and robustness of the target scenarios compared to the equipment and network configurations that were used for this study so far.

In general the observed performance of the tested representative network configurations can support 3GPP use cases with the sole exception of the most stringent motion control case. Other sub-categories of motion control require 3 and 6 msec round trip performance which can readily be achieved with proper network planning and equipment selection. For these use cases networks should be limited to the 1-hop configuration to ensure that the required performance is achieved.

Multiple ingress/egress points are critical for fault tolerance and for system performance as well. Mesh is best suited for this, but a nearly equivalent situation can be achieved with multiple Point-to\_Multi-Point (PtMP) deployments provided the clients can reassociate as required to meet the system performance criteria. IA networks should be designed for service delivery aspects, but this can lead to complications if reconfigurations are needed. Packet sizes and repetition rates are key parameters for IA network designs.

As observed in our testing some equipment can exhibit uncontrolled behaviors. This may limit the addressable target applications in the short-term, but with new chipsets and equipment (e.g. 802.11ay) these limitations are likely to disappear. Another aspect of 802.11ay that could be very beneficial in this regard is the TDD scheduling feature which should regulate the traffic flows through the network much more than the listen-before-talk behavior of 802.11ad.

In conclusion, 60 GHz networks can deliver < 2 msec round-trip time (RTT) latency and meet 3GPP use case performance requirements today with proper equipment selection and network design for substantially less total cost of ownership (TCO) and will offer better capabilities for addressing these use cases in the near future.

Acknowledgements: The author would like to thank Pooja Shankar for conducting much of the testing done in the first indoor trial.

## Abbreviations

AP	Access Point
DUT	Device Under Test
GHz	GigaHertz
HMI	Human-Machine Interface
Hz	Hertz
IA	Industrial Automation
IT	Information Technology
LBT	Listen Before Talk
Mbps	Megabits per second
MEC	Multi-access Edge Computing
MIMO	Multiple Input Multiple Output
msec	millisecond
MSO	Multiple-System Operator
PtMP	Point to Multi-Point
RTT	Round Trip Time
RU	Remote Unit
TDD	Time Domain Duplexing
UE	User Equipment
UPF	User Plane Function
URLLC	Ultra-Reliable Low-Latency Communication
3GPP	3 <sup>rd</sup> Generation Partnership Project
5GC	5G Core

## Bibliography & References

- [1] Instructor comments from *Implementing Industry 4.0: Leading Change in Manufacturing and Operations*, MIT Sloan Executive Education, July, 2019
- [2] Course materials from *Implementing Industry 4.0: Leading Change in Manufacturing and Operations*, MIT Sloan Executive Education, July, 2019
- [3] 3GPP TS 22.104: *Service requirements for cyber-physical control applications in vertical domains*, Annex A
- [4] 3GPP TS 22.261: *Service requirements for the 5G system*, Annex C
- [5] IEEE 802.11r-2008 has been folded in as Section 13 of IEEE 802.11-2016
- [6] *The Digital Shopfloor - Industrial Automation in the Industry 4.0 Era: Performance Analysis and Applications*, John Soldatos et al (editors), River Publishers, May, 2019
- [7] Network Computing, *WiFi Fast Roaming, Simplified*, <https://www.networkcomputing.com/wireless-infrastructure/wifi-fast-roaming-simplified>, retrieved August 6<sup>th</sup>, 2020
- [8] Chen et al, *Millimeter-Wave Fixed Wireless Access Using IEEE 802.11ay*, IEEE Communications Magazine, Vol. 57, Issue 12, December 2019

[9] RCR Wireless News, *Bringing 5G NR to unlicensed spectrum*, <https://www.rcrwireless.com/20181102/5g/5g-nr-unlicensed-spectrum>, retrieved August 13<sup>th</sup>, 2020

[10] Research and Markets, *Global Millimeter Wave Technology Market Size, Market Share, Application Analysis, Regional Outlook, Growth Trends, Key Players, Competitive Strategies and Forecasts, 2019 To 2027*, March, 2020

[11] Fortune Business Insights, *5G Infrastructure Market Size, Share and Industry Analysis By Component (Fibers, Cables, Antenna, Transceiver, Wireless Backhaul, Modem, Router), By Communication Infrastructure (Small Cell, Macro Cell, Radio Access Network (RAN), Distributed Antenna System (DAS)), and Regional Forecast 2019-2026*, July, 2019

# **50 Million Keys to SNMPv3 Privacy**

A Technical Paper prepared for SCTE•ISBE by

**Paul E. Schauer**

Distinguished Engineer

Comcast Cable

183 Inverness Drive West, Englewood, CO, 80111, USA

+1-303-372-1215

paul\_schauer@comcast.com

# Table of Contents

Title	Page Number
1. Introduction.....	3
2. Don't Panic .....	3
3. Get On With It.....	4
4. Cue the Diffie Hellman Key Exchange.....	4
4.1.    SNMPv2 Community.....	4
4.2.    SNMPv3 User-based Security Model with no authentication or privacy.....	4
4.3.    SNMPv3 User-based Security Model with authentication and privacy .....	5
4.4.    SNMPv3 USM authentication and privacy with DH key exchange .....	5
5. Roll a Random Number.....	5
6. Keeper of the Keys.....	6
7. Pick User Profiles .....	7
8. Break Out the Bootfiles .....	7
9. Double Click .....	7
10. Prime the Pollers .....	8
11. Conclusion.....	9
Abbreviations .....	9
Bibliography & References.....	9
Appendix 1 RFC 2786 DH key exchange simulation .....	10

## List of Figures

Title	Page Number
Figure 1 - SNMPv2 Community Example .....	4
Figure 2 - SNMPv3 No Authentication No Privacy Example .....	5
Figure 3 - SNMPv3 Authentication + Privacy Example.....	5
Figure 4 - SNMPv3 Authentication + Privacy + Local Keys Example.....	5
Figure 5 - CM Public Key Query Example .....	8

## 1. Introduction

Security and Privacy must be top line features of any service operating today. Twenty years ago, the team at CableLabs prescribed SNMPv3 with Diffie Hellman key exchange for encrypting management traffic from DOCSIS cable modems. While SNMP is dated, simple non-sensitive management data is still integral to CM operations. Yet in 2020, few production implementations of encrypting SNMP have been reported. Because CM management traffic lives behind numerous other layers of network and physical security, this has not been a significant issue. Still, leaving the CM management data in clear text SNMP is an operational luxury that should be phased out of practice. Switching to newer, more secure protocols is a future solution that overlooks the millions of installed DOCSIS CMs. As part of Comcast's ongoing evolution of security and privacy, clear text SNMP data has been deprecated. DOCSIS cable modem management is now encrypted with SNMPv3 utilizing Diffie Hellman key exchange as specified by CableLabs. Associated mechanisms for securely managing the privacy keys and bootfiles are part of the larger solution. This paper will highlight technical issues of implementing SNMPv3 with Diffie Hellman key exchange at MSO scale.

## 2. Don't Panic

The information exchanged over SNMP is focused on operating network equipment. The information includes data such as the uptime of the device, the list of interfaces and their name, and the number of packets going in and out of the device. There should not be any sensitive material such as credit card numbers, account numbers, or other identifiers contained in the data. Since the first two versions of the protocol had no security, the optimistic expectation for implementors was "don't put sensitive data into SNMP MIBs".

This understanding may have been adequate in the 1980s, when SNMP was first created. In 2020, this is no longer tenable. Various SNMP-based denial of service attacks pose new threats. In addition, privacy requirements that have evolved make this a protocol that should be removed from modern networks.

There are hundreds of millions of DOCSIS cable modems providing service in cable operators' networks that use SNMP. Replacing all of them to remove SNMP is operationally unrealistic.

This is where adding encrypted SNMPv3 with Diffie Hellman key exchange provides an important layer of privacy for these devices. CableLabs and the IETF released RFC 2786 in 2000, and DOCSIS cable modems are required by the CableLabs DOCSIS CM-SP-OSSI specification to implement the RFC.

Actually deploying DH key exchange at scale was left to implementers: "The configuration [of the cable modem SNMP agent] would be done either manually (in the case of a small number of devices), or via some sort of distributed configuration file. The actual mechanism is outside the scope of this document [RFC 2786]."

This document provides a guide for that out of scope effort. This document does not reveal Comcast's implementation of an SNMPv3 DH key management system. There are decision points and examples for needed subsystems that can constitute an overall solution. Any given implementation will be unique to the environment where it is deployed.

All examples are for illustration purposes only and should not be used in actual production without formal review.

### 3. Get On With It

Key design decision: Get over the debate and get on with the deployment of DOCSIS CM SNMPv3 with DH key exchange.

The MD5 hashing algorithm and the DES encryption algorithm with 56-bit keys prescribed in RFC 2786 have been deprecated. They are obsolete because they can be compromised. The SNMP User Security Model described in RFC 3414 highlights the limitations of the overall SNMP security model.

If the information from the cable modems has a low security value, and it is protected by many other layers of security, why even bother obscuring it with known weak encryption?

This question may be argued numerous times and numerous ways by numerous engineers. They will raise valid technical points. Undertaking the journey to SNMPv3 with DH key exchange starts with managerial direction, not technical requirements. For instance, there may be a corporate directive to remove SNMPv2 in the network. A corporate information policy may require that systems encrypt all data in transit. The key here is to stop the debate and get on with implementation.

### 4. Cue the Diffie Hellman Key Exchange

Key design decision: Validate whether your code base supports the use of local keys for the SNMP agent on the cable modem.

As of this writing, the Net-SNMP utilities and Java's SNMP4j library support the use of local keys. This section summarizes differences between SNMPv2 community, SNMPv3 USM, and SNMPv3 USM with DH key exchange. It uses the Net-SNMP command line tools to demonstrate these differences in the calls to each service. This knowledge is required to plan necessary upgrades to the current SNMP polling software. An SNMP poller is formally an "SNMP Command Generator" or "SNMP manager" and can be implemented in numerous forms.

#### 4.1. SNMPv2 Community

Reinforcing that "S stands for Simple", the SNMPv2 community string is a plain text way to group SNMP managers' and agents' commands and responses. The community string is sent in clear text as part of an SNMPv1 or v2 packet. This offers no security at all nor was it intended to.

```
$ snmpget -v 2c -c commString 10.168.6.82 system.sysUpTime.0
10.53.115.71.58645 > 10.168.6.82.161: { SNMPv2c C="commString"
{ GetRequest(28) R=445233587 .1.3.6.1.2.1.1.3.0 } }
10.168.6.82.161 > 10.53.115.71.58645: { SNMPv2c C="commString"
{ GetResponse(32) R=445233587 .1.3.6.1.2.1.1.3.0=182774513 } }
```

**Figure 1 - SNMPv2 Community Example**

#### 4.2. SNMPv3 User-based Security Model with no authentication or privacy

This may seem counter-intuitive, but SNMPv3 provides a method for unsecured communication. The SNMPv2 "community" evolves into the SNMPv3 "user". There is a default user with default views and



permissions prescribed in RFC 2786 that all DOCSIS cable modems implement. This becomes important later on for exchanging keys. This example uses the default DOCSIS user “dhKickstart” with no authentication or privacy keys in the request.

```
$ snmpget -v 3 -l noAuthNoPriv -u dhKickstart 10.168.6.82
system.sysUpTime.0
```

**Figure 2 - SNMPv3 No Authentication No Privacy Example**

### **4.3. SNMPv3 User-based Security Model with authentication and privacy**

SNMPv3 adds methods for authenticating message transmission and reception. It also provides methods for encrypting the contents of the message in transit. The standard RFC 3414 implementation uses an authentication key and a privacy key that are loaded into the SNMP manager and the SNMP agent(s). The manager and agents use these keys in the algorithms prescribed by RFC 3414 to transmit and decrypt messages in lieu of plain text SNMPv2.

```
$ snmpget -v 3 -l authPriv -u v3keysTest -A TestAuthenticationKey -X
TestPrivacyKey 10.168.6.21 system.sysUpTime.0
```

**Figure 3 - SNMPv3 Authentication + Privacy Example**

### **4.4. SNMPv3 USM authentication and privacy with DH key exchange**

The addition of DH key exchange prescribes a method where the authentication and privacy keys are derived based on cryptographic methods. This improves on the standard SNMPv3 implementation because it removes the actual keys from the cable modem configuration. This example shows authentication and privacy keys that are precomputed based on the DH key exchange, and are then used in the call to the Net-SNMP utilities.

```
$ snmpget -v 3 -l authPriv -u v3localKeys -3k
0xf9c96a9232ee65a08aa9085e6f1ff82c -3K
0xedebac85112645218fb7a63659a332c2 10.168.6.21 system.sysUpTime.0
```

**Figure 4 - SNMPv3 Authentication + Privacy + Local Keys Example**

These examples are simple demonstrations of how calls from an SNMP manager to an SNMP agent change. The specific code used in a cable provider’s system will need to be upgraded similarly.

## **5. Roll a Random Number**

Key design decision: Creating and managing the distribution of the Manager Random number(s).

The manager random number is the basis for all subsequent activities associated with RFC 2786 DH key exchange. A random number is transformed according to the algorithms in RFC 2786 to arrive at the

manager public key that must be included in the cable modem boot file. Each calculated public key is exposed through the CM's SNMP MIB as a "usmDHKkickstartMgrPublic" key. Each SNMP user profile must have its own manager random number to create its unique key. The cable modem follows a similar process and creates its own "usmDHKkickstartMyPublic" key for each user as well. These values are then further processed and used in calls from the SNMP manager to the SNMP agent as shown in the previous section.

The simplest and least secure way to manage the random number is to make it available to the SNMP polling systems the same way the SNMPv2 community string was provided. Methods such as email, text files, or hard coded configuration pushes allow any system to calculate the DH keys based on the manager random number. This is also the least secure method a system designer could choose.

More secure solutions also require more effort to manage. Code repositories or secured APIs with token authentication schemes can be used. A formal key vault or similar key management infrastructure may be used to authorize access to a given manager random number for a given user.

Since an SNMP manager needs to be able to calculate the authentication and privacy keys any time a cable modem reboots, the scaling of the system which manages needs to be carefully considered and expanded to accommodate the potential query traffic. The parameters a system designer needs to consider include: the number of SNMP USM profiles with unique manager random numbers, the number of different boot files in the system, and the frequency of change of the manager random numbers. Since an SNMP manager could calculate the keys itself, the number of cable modem reboots in the overall system is less important at the manager random number distribution level.

The system designer may also wish to avoid handing out the manager random number(s) completely, and instead offer a service that provides the pre-calculated keys to requestors. This keeps the manager random numbers from being cached or redistributed to other parties.

## **6. Keeper of the Keys**

Key design decision: Manage the distribution of calculated USM keys in lieu of sending out manager random numbers

Another layer of flexibility can be added with a USM key management service. Instead of directly sending out the manager random numbers, a dedicated service can be created to produce the authentication and privacy keys to requestors.

There are security benefits to a USM key distribution service. The manager random numbers are not divulged to the requesting SNMP manager, only the calculated keys. This improves the security of the overall system. It also removes the necessity of SNMP managers to calculate the keys themselves, which may require significant code updates.

There are also drawbacks to distributing calculated keys. The primary challenge is scalability. Since the cable modem public keys change every time the cable modem reboots, the service must be able to refresh the CM public key from any modem when it reboots. A separate key distribution also creates a serial failure scenario. If SNMP managers must retrieve the precomputed keys from a key distribution service, and the service is impaired or down, visibility will be lost to the CMs as they reboot over time.

A separate USM key service like this may be used for all the users and all the keys, or it may be selective for only the user(s) that have read and write access. This may provide a compromise between potential loss of visibility, query rate, and key integrity.

The final consideration for a USM key service is the potential query rate. The CM keys change every time they reboot, hence the USM service must retrieve and recalculate the keys each time. Then the new keys must be made available to any SNMP manager using the service. The number of cable modems, the number of reboots, the number of SNMP managers in the system, and the frequency of SNMP queries from the SNMP manager are all scaling factors that will result in the potential query rate the key store would need to accommodate.

## **7. Pick User Profiles**

Key design decision: SNMP users and roles to include in the CM bootfile to generate DH keys. Each user profile will create its own public key based on its own manager random number, which will then be used to calculate the local keys for that user.

The unsecured “dhKickstart” user documented in RFC 2786 has a basic view and access to the public keys necessary to calculate the actual CM local keys. Additional users should be added with appropriate views and permissions to monitor the cable modems. Operators may leverage existing RFC 3414 USM and RFC 3415 VACM profiles in their cable modem boot files. These may be copied or further refined as desired for SNMPv3 with DH key exchange.

There is nothing extraordinary about these users or profiles. The one decision point of note is whether to partition a read/write user and read/only users.

## **8. Break Out the Bootfiles**

Key design decision: How the system will manage the bootfiles with the new users and keys.

Cable modem provisioning systems already manage multiple bootfiles. Adding the SNMPv3 users and manager random numbers to the bootfiles may increase this quantity. There are numerous ways for a cable operator to partition the bootfiles. Groupings might include the CM manufacturer, model number, physical location, organizational responsibility, provisioning complex, or other criteria.

Dynamically generated bootfiles with a unique manager random number for each user on each cable modem would be the most secure solution. If an actor obtained one manager random number for a given CM, an SNMP session could be negotiated to only that cable modem with the specific user. The manager random number would change when that CM rebooted, and all the security parameters would be reset. This is also the most difficult scenario to manage. The provisioning system that generates the bootfiles would also have to seed the distribution system for all the manager random number(s) as described above.

A single random manager number that is used for the keys in every bootfile is the easiest solution, and clearly the least secure. Any actor that obtains the manager random number could securely negotiate DH sessions with any cable modem in that system.

Deciding how many users and manager random numbers could thus simply leverage the existing bootfile management method. It could also completely change the management and distribution of them.

## **9. Double Click**

Key design decision: Whether SNMP pollers rely on cached keys, or query and calculate keys for each session.

As shown below, querying the usmDHKickstartMyPublic key(s) from a given cable modem with the default dhKickstart user should always return the cable modem public keys.

```
$ snmpbulkwalk -v 3 -m ALL -l noAuthNoPriv -u dhKickstart 10.168.6.82.161
snmpUsmDHObjectsMIB
SNMP-USM-DH-OBJECTS-MIB::usmDHKickstartMyPublic.1 = Hex-STRING:
97 BC 5A C7 B8 FC 32 80 07 2E 20 4A CE 6D 59 86
0E 0B C8 B2 F7 DD 86 4F 7D B3 E6 21 B2 51 99 32
FE 62 9A 20 B3 CD 72 97 5A 18 B5 E1 87 02 89 AB
6B 67 9F 38 1C BA E4 07 A8 BA 9D 80 40 BE 3B 26
C6 B9 9E F8 D3 70 0E A3 3A 34 95 3F 51 A2 55 95
EC AE FF 84 A9 72 57 8C AB 36 45 76 8B 7F 32 95
C6 BD 93 D9 DB CF 6A 2E 05 10 CE 9E 2E 4A 01 CB
1B 27 42 68 25 BD 54 55 79 79 93 AE 8C 61 47 AB
```

**Figure 5 - CM Public Key Query Example**

Whether the SNMP manager is a subsystem within a USM key management service described previously or a separate manager that has access to the SNMP manager random number for a given user, the fact that the keys only change when the cable modem reboots can be used as a convenience to avoid double querying the modems.

## 10. Prime the Pollers

Key design decision: Scaling the SNMP poller inventory to accommodate the key exchange and encryption overhead.

The additional response time from the CMs with the key exchange and encryption at MSO scale is non-trivial yet manageable. Response time measurements of agent response in laboratory settings have shown small increases of a few milliseconds or less. Scaling out to CMTSs with thousands of CMs and carrying full customer traffic have shown increases from a few milliseconds to a few seconds depending on CM and CMTS load at the time. Some phenomena observed in the production network include:

“Warm up failures”, or CM key exchanges that take longer than expected.

Serialized key negotiation to sets of modems that should be converted to parallel queries.

Data response times that take randomly longer than the SNMPv2 equivalent.

The solution for a system designer is the same: additional SNMP pollers with code optimized for parallelism. Experience at Comcast’s scale shows a 30% to 300% increase in resource needs depending on the polling software and computing platform. Physical servers, virtual containers, threading models, cloud computing instances, and more may need to be scaled specifically to meet the overall system demands. There has not been a universal solution observed at Comcast’s scale; all of the above have needed some adjustment based on full production deployment. The good news is they were all identified during the roll out and addressed so that the project was deployed successfully.

## 11. Conclusion

Adding privacy by encrypting cable modem management traffic starts by opening the door to RFC 2786. The keys in this presentation are provided to help guide design decisions necessary to implement SNMPv3 with DH key exchange on CMs at production scale. The balance between maintainability and scalability of each key must be considered as part of the overall system design and implementation.

## Abbreviations

CM	cable modem
DH	Diffie Hellman
DOCSIS	data over cable service interface specification
PBKD	password based key derivation
RFC	request for comments
SNMP	simple network management protocol
SCTE	Society of Cable Telecommunications Engineers
USM	user security model
VACM	view-based access control module

## Bibliography & References

RFC 2786: *Diffie-Hellman USM Key Management Information Base and Textual Convention*; RFC Editor

RFC 3414: *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*; RFC Editor

RFC 3415: *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*; RFC Editor

# Appendix 1 RFC 2786 DH key exchange simulation

All of the constants and mathematics associated with the DH key exchange are documented in RFC 2786. The following illustrates these calculations using open source tools. The methods shown are long, single steps for demonstration purposes. Actual systems implementation would be succinct.

Pick a 1024 bit manager random number:

```
$ python3 -c 'import secrets; print(hex (secrets.randbits(1024)))'
```

```
0x8ad7112ed1742765da233bb761f44cc69b329161e1a2a1c28032c25445566284bf6ca06a8a91b2a49c06
b7198aeb1574570b1e1ab22f7a9c471d01ed1f260f7356e3c88b92f06e70a14ac91480677aa4c571fb0d1c
f0a9857d027f98b1426701ab6fb27a933b0db3821b7c401fdbcb7f4077ae87c4f64f682a0ffe54ff5c90928
```

Then transform the random number as prescribed using the prime number listed in the RFC The equation to derive the keys is:  $((\text{base number})^{\text{random number}}) \text{ modulo DH prime number}$  ):

```
$ python3 -c 'print(hex(pow (2,
0x8ad7112ed1742765da233bb761f44cc69b329161e1a2a1c28032c25445566284bf6ca06a8a91b2a49c06
b7198aeb1574570b1e1ab22f7a9c471d01ed1f260f7356e3c88b92f06e70a14ac91480677aa4c571fb0d1c
f0a9857d027f98b1426701ab6fb27a933b0db3821b7c401fdbcb7f4077ae87c4f64f682a0ffe54ff5c90928
),
0xffffffffffffffffc90fdaa22168c234c4c6628b80dc1cd129024e088a67cc74020bbea63b139b22514a
08798e3404ddef9519b3cd3a431b302b0a6df25f14374fe1356d6d51c245e485b576625e7ec6f44c42e9a6
37ed6b0bff5cb6f406b7edee386bfb5a899fa5ae9f24117c4b1fe649286651ece65381fffffffffffffffff
)))'
```

The result is a “usmDHKickstartMgrPublic” number that would then be included in the CM bootfile and exposed by the CM SNMP MIB:

```
0xef26516697b54859842f6fdc933a9aeac3750a1d50472808e62f845efb604caf5e97b117b061538a057b
1fef375da865cc0344d36f7cc6e0cc7265aa439f423d250c20518ba8459df50d710ba30a4292d43b7fbcee
189faa12cafdafd54c4797aac2bdb5438b5f7a8a5df28b12673b95a8a17682361815fd026960bb5ab021dc
```

Internally, the cable modem follows the same algorithm each time it reboots:

```
$ python3 -c 'import secrets; print (hex (secrets.randbits(1024)))'
```

```
0x94815e16cb9ac643050278ef80825a384f946d3c050228c440700a7cfd7b4f4acc58b095d51be81ec11e
d5129434588384e771d98d327282fc9d4ff1b81930563a0d5f988952f6b166cbcd6bb692d8c1b2c320762d
2b24d5eea0a0fde537424f2fc09f8a1d65539d26472bc45b60bbc2ea03d2fa98b996c8677e0419c4fe825d
```

```
$ python3 -c 'print (hex (pow (2,
0x94815e16cb9ac643050278ef80825a384f946d3c050228c440700a7cfd7b4f4acc58b095d51be81ec11e
d5129434588384e771d98d327282fc9d4ff1b81930563a0d5f988952f6b166cbcd6bb692d8c1b2c320762d
2b24d5eea0a0fde537424f2fc09f8a1d65539d26472bc45b60bbc2ea03d2fa98b996c8677e0419c4fe825d
),
0xffffffffffffffffc90fdaa22168c234c4c6628b80dc1cd129024e088a67cc74020bbea63b139b22514a
08798e3404ddef9519b3cd3a431b302b0a6df25f14374fe1356d6d51c245e485b576625e7ec6f44c42e9a6
37ed6b0bff5cb6f406b7edee386bfb5a899fa5ae9f24117c4b1fe649286651ece65381fffffffffffffffff
)))'
```

The result is a simulated “usmDHKickstartMyPublic” key that would be exposed by the CM:

```
0x992d538d0cf24e4869b3cdc8e04f7ee6c6fb454b2b7b412a585dc09d9c293167cea9af58710695ff9344
ca7770da876f3124034af013933d257cddc21930886b364dcad417d8af56ed8b7b6953e7de5f8f0ad3c5a1
5f85578d9a64a7900bb58646c8bf804dc099da5428d1308db33c4cab828ca1618e2258561ba66f9c23ad3
```

Now that both the CM and the SNMP manager have known public keys, each calculates a shared secret using the other's public key:

At the SNMP manager:

```
pow(usmDHKickstartMyPublic, manager random number, DH prime)
```

```
$ python3 -c 'print (hex (pow
(0x992d538d0cf24e4869b3cdc8e04f7ee6c6fb454b2b7b412a585dc09d9c293167cea9af58710695ff93
44ca7770da876f3124034af013933d257cddc21930886b364dcad417d8af56ed8b7b6953e7de5f8f0ad3c5
a15f85578d9a64a7900bb58646c8bf804dc099da5428d1308db33c4cab828ca1618e2258561ba66f9c23ad
3,
0x8ad7112ed1742765da233bb761f44cc69b329161e1a2a1c28032c25445566284bf6ca06a8a91b2a49c06
b7198aeb1574570b1e1ab22f7a9c471d01ed1f260f7356e3c88b92f06e70a14ac91480677aa4c571fb0d1c
f0a9857d027f98b1426701ab6fb27a933b0db3821b7c401fdcb7f4077ae87c4f64f682a0ffe54ff5c90928
),
0xFFFFFFFFFFFFFFFFC90FDA22168C234C4C6628B80DC1CD129024E088A67CC74020BBEA63B139B22514A
08798E3404DDEF9519B3CD3A431B302B0A6DF25F14374FE1356D6D51C245E485B576625E7EC6F44C42E9A6
37ED6B0BFF5CB6F406B7EDEC386BFB5A899FA5AE9F24117C4B1FE649286651ECE65381FFFFFFFFFFFFFFFF
)))'
```

Shared secret at the SNMP manager:

```
0xf4729f7ec7f55e997060e85688b72efa84fb0e1d2aa29bc95f07a5cfc4e2542957147144a43ddba41c66
b0372985b42f3222da171340a4737e0d7a62643b2894aa27c9c66a67f95455b84f3f5226d76adc6bdf9893
a3149eb0529460651df62b0afec21370a9e92eff84062b12878ff5b19fb8029f6d089d2c7fe8576ff694b
```

Simulating the internal process at the Cable Modem:

```
pow(usmDHKickstartMgrPublic, CM random number, DH prime)
```

```
$ python3 -c 'print (hex (pow
(0xef26516697b54859842f6fdc933a9aeac3750a1d50472808e62f845efb604caf5e97b117b061538a05
7b1fef375da865cc0344d36f7cc6e0cc7265aa439f423d250c20518ba8459df50d710ba30a4292d43b7fbc
ee189faa12cafdaafd54c4797aac2bdb5438b5f7a8a5df28b12673b95a8a17682361815fd026960bb5ab021
dc,
0x94815e16cb9ac643050278ef80825a384f946d3c050228c440700a7cfd7b4f4acc58b095d51be81ec11e
d5129434588384e771d98d327282fc9d4ff1b81930563a0d5f988952f6b166cbcd6bb692d8c1b2c320762d
2b24d5eea0a0fde537424f2fc09f8a1d65539d26472bc45b60bbc2ea03d2fa98b996c8677e0419c4fe825d
),
0xFFFFFFFFFFFFFFFFC90FDA22168C234C4C6628B80DC1CD129024E088A67CC74020BBEA63B139B22514A
08798E3404DDEF9519B3CD3A431B302B0A6DF25F14374FE1356D6D51C245E485B576625E7EC6F44C42E9A6
37ED6B0BFF5CB6F406B7EDEC386BFB5A899FA5AE9F24117C4B1FE649286651ECE65381FFFFFFFFFFFFFFFF
)))'
```

Shared secret at the CM, which is exactly the same as the calculated shared secret at the SNMP manager:

```
0xf4729f7ec7f55e997060e85688b72efa84fb0e1d2aa29bc95f07a5cfc4e2542957147144a43ddba41c66
b0372985b42f3222da171340a4737e0d7a62643b2894aa27c9c66a67f95455b84f3f5226d76adc6bdf9893
a3149eb0529460651df62b0afec21370a9e92eff84062b12878ff5b19fb8029f6d089d2c7fe8576ff694b
```

The derivation of the authentication and privacy keys then follows the method in RFC 2786. There are several default parameters specified in the RFC that are used for the derivation using the PBKDF (password based key derivation) algorithm.

Authentication salt value: 0x98dfb5ac

Privacy salt value: 0xd310ba6

Both keys use 16 bit lengths and 500 iterations of the algorithm.

Authentication key:

```
python3 -c 'import pbkdf2; import binascii; print ("0x" + binascii.hexlify
(pbkdf2.PBKDF2
(0xf4729f7ec7f55e997060e85688b72efa84fb0e1d2aa29bc95f07a5cfc4e2542957147144a43ddba41c
66b0372985b42f3222da171340a4737e0d7a62643b2894aa27c9c66a67f95455b84f3f5226d76adc6bdf98
93a3149eb0529460651df62b0afec21370a9e92eff84062b12878ff5b19fb8029f6d089d2c7fe8576ff694
b.to_bytes(128, "big"), 0x93dfb5ac.to_bytes(4, "big"),
iterations=500).read(16)).decode()) '
```

0xbe0e87f5a70e2ae6ad12a513ea029c16

Privacy key:

```
python3 -c 'import pbkdf2; import binascii; print ("0x" + binascii.hexlify
(pbkdf2.PBKDF2
(0xf4729f7ec7f55e997060e85688b72efa84fb0e1d2aa29bc95f07a5cfc4e2542957147144a43ddba41c
66b0372985b42f3222da171340a4737e0d7a62643b2894aa27c9c66a67f95455b84f3f5226d76adc6bdf98
93a3149eb0529460651df62b0afec21370a9e92eff84062b12878ff5b19fb8029f6d089d2c7fe8576ff694
b.to_bytes(128, "big"), 0xd310ba6.to_bytes(4, "big"),
iterations=500).read(16)).decode()) '
```

0x50ed4b963e81ff58227497a966bacd6c

These two calculated keys would then be used by the SNMP manager to communicate with the cable modem SNMP agent. When the CM reboots, or the manager random number changes, the keys must be recalculated.



# Dynamic Data Collection & Configuration Management

A Technical Paper prepared for SCTE•ISBE by

**Rohini Vugumudi**

Director, Software Development, Comcast  
Comcast  
Rohini\_vugumudi@comcast.com

Co-Authors:

**Hany Fame**

Software Developer 5, Comcast  
Hany\_Fame@cable.comcast.com

**Pardeep Singh**

Software Developer 5, Comcast  
Pardeep\_Singh@cable.comcast.com

**Zhen Lu**

Sr. Software Developer, Comcast  
Zhen\_lu2@comcast.com

# Table of Contents

Title	Page Number
1. Introduction.....	3
2. Historical Context .....	3
3. The Solution .....	5
3.1 Data Collection Platform (Genome) .....	6
3.2 Configuration Manager.....	10
3.2.1 Configuration Manager State Machine Scheduler .....	11
3.2.2 CLI Commands Generation and CMTS Device Configuration .....	11
3.2.3 Configuration Manager State Machine Supports and Uses the Following Technologies:...	13
4. Challenges / Discoveries.....	14
4.2 Genome.....	14
4.3 Configuration Manager.....	15
5. The Future .....	17
6. Conclusion.....	20
Abbreviations .....	20
Bibliography & References.....	20

## 1. Introduction

The modulation profiles for Cable Modem Termination Systems (CMTSs) have been historically applied manually, only changing with response to stimuli such as frequency impairments identified by field engineers or, worse, customers. This manual feedback loop is inherently slow, resulting in profile configurations that are limited by the impairments of the lowest common denominator on a given cluster of customers. With the advent of Data Over Cable Service Interface Specification (DOCSIS) 3.1 came the amazing and powerful ability to automatically adapt downstream profiles in near real-time leveraging machine learning and the Profile Management Application (PMA) concept. Tightening and automating the feedback loop allows for the recovery of previously wasted capacity, thereby making the entire network more efficient.

The authors of this publication have developed an implementation of the PMA concept that allows Comcast to manage the downstream and upstream environments efficiently at scale. Our current nascent architecture allows us to run a complete feedback loop every 6 hours; this runtime should only get better with further optimization of the Analytic Engine (AE) service, which is currently the bottleneck. Comcast can now optimize for the best performance, reliability and throughput in an automated and scalable fashion.

In this paper, the service architecture platform for dynamic data collection is called Genome, and it offers a key component of the overall configuration management system. Genome is responsible for the aggregation of data collected from cable modems and other customer devices, and the configuration management service is responsible for the application and validation of generated modulation profiles to their respective parent CMTSs. In particular, the details of adapting modern cloud computing tools to architect a reliable software solution for both downstream and upstream configurations are discussed. There are many details involved in the data aggregation, application of configurations, and validation of configurations, all of which are discussed.

## 2. Historical Context

The ability for cable providers to manage communications parameters for essential hardware has improved dramatically over the past several decades. The industry began with hardware that was statically configured or configured by physical manipulation. The invention of the DOCSIS specification allowed for configurable modulation profiles; however, profiles were statically defined by nature. In addition, while there is physically no difference between downstream and upstream channels, they have evolved differently over time due to how they have been historically allocated and used. This section will cover the evolution of both channel types and the difficulties in data collection and configuration management that arose because of this.

### **Data collection Limitations:**

It is not a trivial task to collect the large amount of data required to create a continuous feedback loop for profile management at the scale of a national footprint. Previous attempts at creating a standard approach have been found wanting in terms of speed, cost, and maintainability. This section covers the limitations of common historical methods of data collection.

The most basic form of data collection has been to allow network engineers to manually record data in spreadsheets. This method is undeniably easy to implement but has many pitfalls. The most obvious is that the minimal feedback loop time is limited by human processing. In order to avoid potentially impairing customer experiences due to configuration changes, the sampling time must be reduced to as short a time frame as possible. Other difficulties that are associated with this method are those of standardization and the inclusion of human error. It is imperative that data sets are well sanitized before they are fed into a machine learning algorithm as input. Last, but not least, most solutions are plagued by scale. To consistently and reliably collect data at scale manually would require many network engineers, therefore increasing costs drastically.

As manual data collection can be ruled out due to practical limitations, the next step is to attempt to implement an automated collection scheme. Such a scheme must be able to scale in several different facets such as number of customer devices, number of CMTSs, and number of data points collected. Data collection from physical CMTSs such as the Arris E6000 converged edge router and consumer cable modems such as the Motorola surfboard SB8200 is limited to polling methods due to neither type of device having native software capable of pushing the required data types. One possible solution to scale is to manually segregate by CMTS or group of cable modems and put each partition onto its own Virtual Machine (VM). This is a practical solution given that the network does not change much over time; however, the demand for data services is growing at an unprecedented pace. Adding more data to feed additional features may result in the need to vertically scale VMs, while adding more CMTSs may require provisioning more VMs. While this is feasible, it would require considerable amounts of manpower in order to manage on a constantly changing edge network.

### **Configuration management Limitations:**

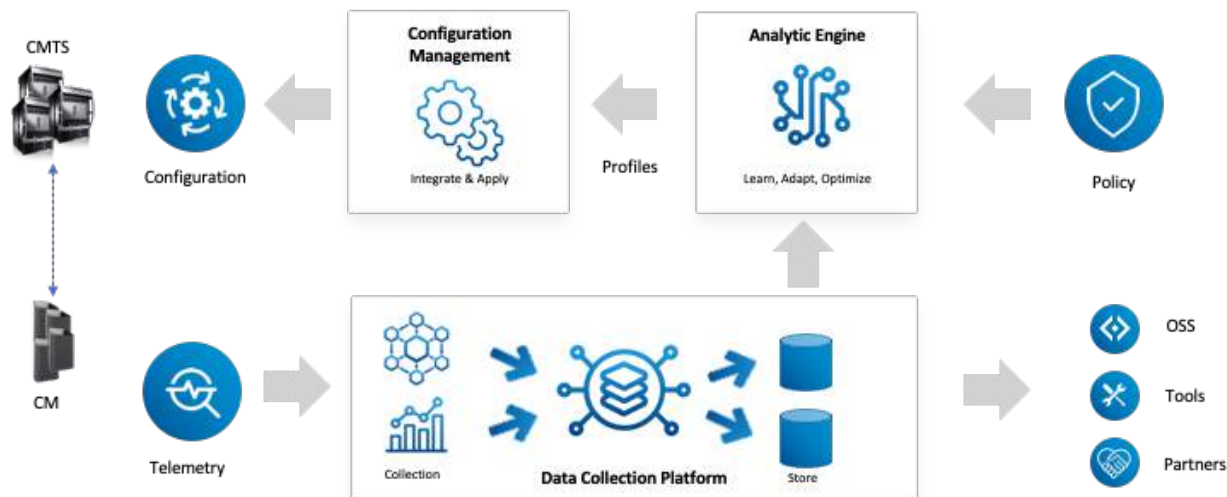
While the problem of dynamic data collection has historically been one of scale, the problem of configuration management has largely been one of standardization, process, performance, and risk. In the recent past, CMTS configurations were statically defined and rarely revisited for optimization due to manpower cost and risk. Even if one were to develop practical methods to collect the required data and use that data to generate optimized configurations, it would still be risky to apply manually said configurations as any mistakes made could negatively impact the customer experience. This fact has resulted in CMTS configurations that are generally very conservative in nature.

The tools and software available for configuration management in the past did not typically yield effective solutions. Most cable companies have, at some point, kept configurations tracked in spreadsheets. As in the case of data gathering, this is prone to human error and is limited in speed due to the human element. Previous solutions have also lacked standardized rules defining how conservative any given configuration should be. As a result, some areas may have more performance-oriented configurations at the cost of possibly degrading the experience of some customers, while others may have tended towards more conservative solutions that may have synthetically decreased the amount of usable bandwidth. Ultimately, the act of changing the configuration for a CMTS has not been something that could be done with the flip of a switch. Numerous validation checks must be completed both to adhere to lawful practices and to ensure that customer service is not impacted negatively. Any such system must also be agile enough to pivot in the case of process failure or external duress.

### 3. The Solution

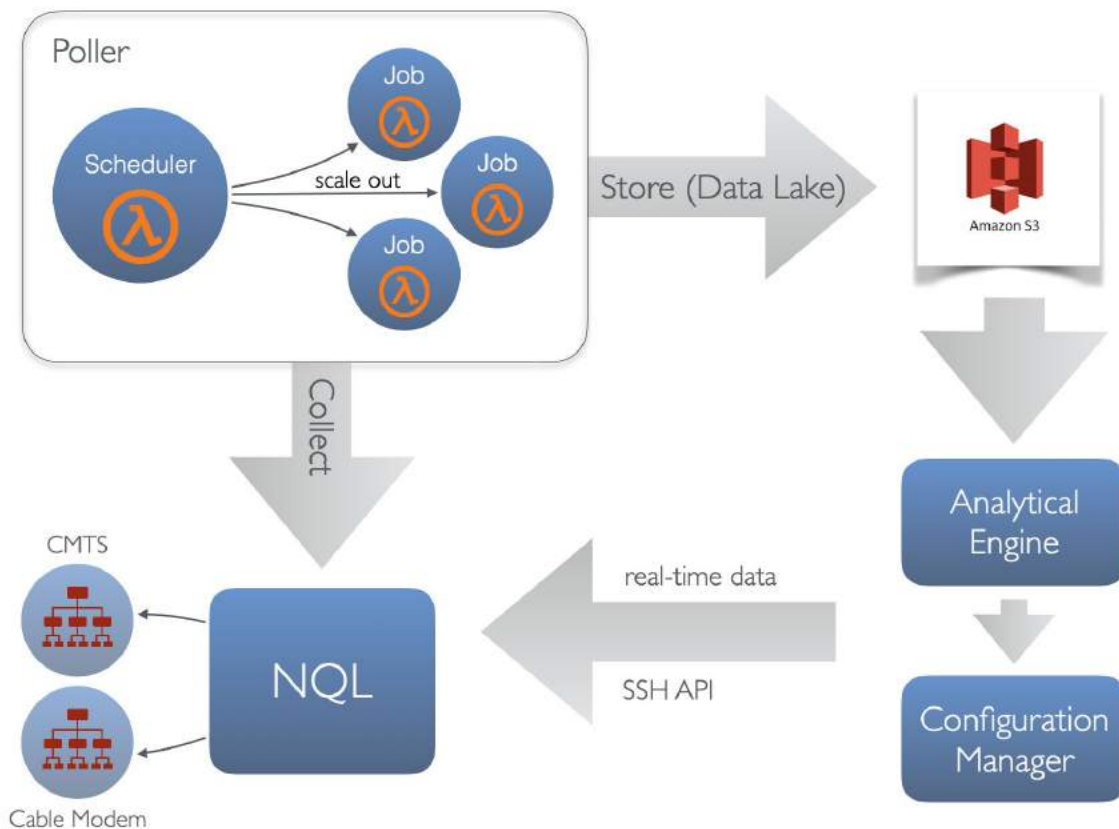
The previous section illustrates the need for a scalable, reliable, and automated approach for the activities involved in dynamic data collection and configuration management. This section discusses one implementation for such an architecture. It is important to note that this architecture is not limited in application to PMA, as it was built to be modular and support other internal processes around operational analysis and field support.

The main driver for the development of these services was to serve Comcast's PMA (Octave) efforts. Considering this, a brief refresher of Octave's architecture is given below before the focus is given to data collection and configuration management.



The first step of the Octave process is data collection, done by the data collection platform, which is internally named Genome. The purpose of this service is to poll both CMTS and cable modem telemetry data, sanitize it, aggregate it, and store it in a data lake for further use. In the case of Octave, the consumer of this data is the Analytic Engine (AE). For the purpose of this publication, this service will be treated as a black box, but interested readers can delve into the details in *A*

*Machine Learning Pipeline for D3.1 Profile Management* Harb 2019. The AE service generates optimized recommendations, then passed to the configuration manager service for translation into configurations. The configuration management service then does all the required validation checks before finally applying the CMTS configurations. This repeats in an infinite loop in order to keep the network configurations optimized.



### 3.1 Data Collection Platform (Genome)

The Genome platform was originally built to serve the Octave initiative but has grown into a platform in its own right since inception. The main purpose of Genome is to actively poll and cache data from devices such as cable modems and CMTSs in a scalable and configurable way while offering consumers the ability to analyze the cached data or get live data in a seamless fashion. This positions Genome to be the sole data provider for cable modem and CMTS data, eliminating inefficiencies around oversaturating CMTSs or cable modems with connections and over fetching duplicative data. In the context of Octave, Genome is responsible for collection, standardization of data required for the creation and management of modulation profiles. Genome

data collection is utilized in the analysis of Proactive Network Maintenance (PNM) activities as well as other field support analysis tools internal to Comcast in addition to Octave.

Genome is made up of two layers, the poller layer, and the query layer. A “poller” is simply a piece of software that is responsible for scheduling, collecting, and standardizing a set of data from devices. Pollers are built to be lean, modular, and extensible. The query layer for Genome is called Network Query Language (NQL), which exposes a declarative API service through which consumers request live or cached data. All pollers are primary consumers of NQL to collect live data, while also offering external consumers to do the same. NQL’s goal is to offer a declarative abstraction layer for edge network devices, allowing consumers to query using standard HTTPS instead of SNMP, TFTP, and various other network communication protocols.

Genome requires a master list of all CMTSs to poll as well as a configuration for the polling cadence for each data property. When a CMTS gets added to the network, it should be added to this master list. In the case of Octave, this master list is automatically updated through Comcast’s deployment ticketing system. On the next polling cycle, Genome will discover the newly added CMTS and request a list of all cable modems it serves. Afterwards, Genome maintains a cached list of all cable modems associated with the CMTS. The list is the source to get the OID data from all saved devices. The polling could be done at any configured frequency or on-demand.

In addition, Genome must ensure that data collection can scale when CMTSs are added over time such that data collection and aggregation can be achieved in a given polling window. It must also ensure that the ingested data is validated and cleaned before it is cached. As the amount of data is large, especially in the case of cable modem data, Genome must also manage data retention policies in order to reduce cost.

The following are some examples of data points that Genome collects:

- OFDM Channel
  - OFDM channel width
  - Subcarrier width
  - Start frequency
  - Active & excluded regions
  - Position of PLC channel
- OFDM Subcarrier
  - Modulation efficiency per subcarrier
  - Subcarrier type
- CMTS
  - Make
  - Model
  - Hardware
  - Software version
- Telemetry
  - MER
  - FEC
  - Traffic

NQL's journey began in the early days of Octave development. The goal was to create an API service which abstracts away all different protocols around networking and let end users interact through HTTPS. The first iteration came in the form of a REST API which accepted OIDs parameters and returned the output through HTTPS. While this iteration ran in productions for several months, several shortcomings were brought to light. For example, use cases involving many data points per CMTS, necessitated making multiple requests to the API. Each request would connect to the CMTS, which resulted in opening and closing sockets many times over multiple requests. Furthermore, the work of encoding and decoding large amounts of JSON data was repeated for each request, degrading the overall performance of the service through repetition of work. It became clear that the architecture needed to evolve in order to overcome these shortcomings. Indeed, with the vision that, Genome, and by extension, NQL, would grow to encompass more and more, it was clear that a technology would have to be chosen to all for incrementally extending the codebase without causing breaking changes at every turn. GraphQL was a natural option that allows data schemas to easily evolve and can handle complex relationships between data sets. NQL was born.

As NQL offers an abstraction layer for connectivity to both cable modems and CMTSs, it must therefore manage all writes as well as reads to each network device. In order to accomplish such a feat, it supports both IPv4 and IPv6 communication protocols. NQL is built primarily using the GraphQL specification and has been optimized at node level. In some cases, consumers may require a request-response type handshake, where the consumer keeps a socket open until the requested data is returned. However, many queries may be long running and it may not be practical or possible for a consumer to keep a connection open for the duration of the process.

Nodejs Historically, the team used Nodejs as a declarative jack-of-all language. However, in this use case, performance is king. As Golang is routinely touted as a performance-oriented language geared towards networking applications, it was a simple choice after orchestrating some internal benchmarking. Simple GraphQL APIs were built in both languages to test a small fraction of our total scale. The goal was to see how many VMs would be needed to handle the same amount of work while profiling the performance of each VM by measuring CPU and Memory usage. As Nodejs is single-threaded language, it is more difficult to saturate a larger VM without adding complexity. To keep it a simple apples to apples comparison, small VMs were used for both languages. After few weeks of tweaks and testing, the conclusion was that Golang is around 3-4 times faster in handling the same work as Nodejs for this use case. This was not a perfect apple to apples comparison since many variables are in play here. Particular to note is the reliance on third party open-source packages, whose performance is out of our control. Another thing that was found is that Golang was very consistent in its response times and overall API metrics, while the API built using Nodejs would suffer spikes in response times. This is likely due to node's single-threaded nature, which makes it not well suited for CPU intensive tasks that block the main execution thread. Golang was the clear winner in this exercise and was used going forward.

NQL's SSH feature allows consumers log into a supported remote device through HTTPS rather than SSH. This allows NQL to abstract away details around authentication, authorization, and managing the underlying SSH connection. Users are able to connect to a host, run multiple commands, and get output back for each command. NQL uses GraphQL to define the API, which



allows developers to develop powerful features and evolve our API without breaking the world. NQLs SSH is a simple layer around secure shell, all the logic around what commands to run in what order and what to do with the output is handled by the client itself. For example, CM can apply configurations to a CMTS and then analyze the output to see if everything went well. All this is done through HTTPS using NQL. NQL has evolved from a standalone service running on a EC2 VM to where it could be leveraged anywhere within our codebase and able to run in any container or lambda alongside our codebase.

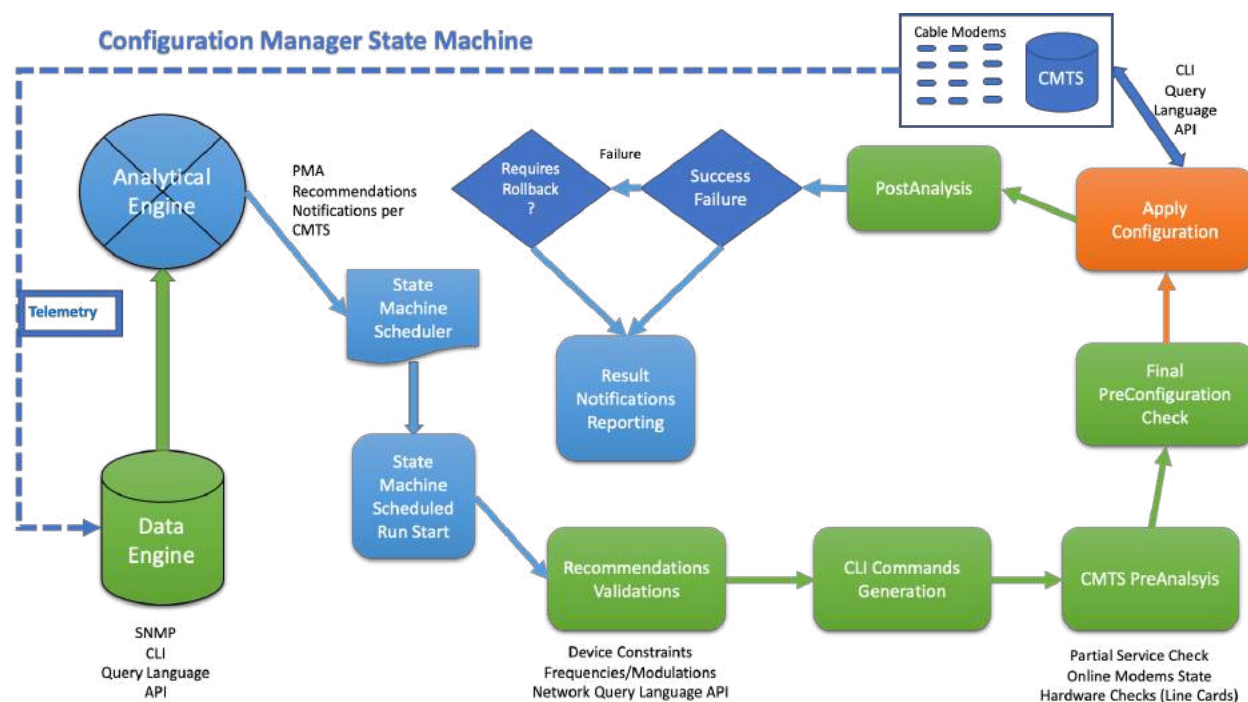
#### SNMP v2/v3 capability

- Genome by default uses the polling on using SNMP V3, and if the device is not V3 enabled, the fall back option is V2.
  - Get public and manager key from cable modem
  - Use manager key and V3 security name as input to get vault secret
  - Use cryptographic hash function on vault secret and public key and compute auth priv keys.
- Collect the SNMP data and stream to Kinesis which will be used by Analytical Engine from PMA and other consumers of the data as need basis.
- AE consumes the data and suggests the profile which will be sent to configuration manager, which further defined below how the process works and how the recommended profile has been applied to using the configuration manager

### 3.2 Configuration Manager

The Configuration Manager (CM) service is a scalable state machine application which is able to generate and apply modulation profiles to CMTSs based on suggestions provided to it by the AE service. It was designed to enhance DOCSIS3.1/3.0 downstream/upstream capacity and correct for RF noise impairments without manual interaction. Unlike the Genome service, it is currently not a standalone service and requires input from the AE service to function.

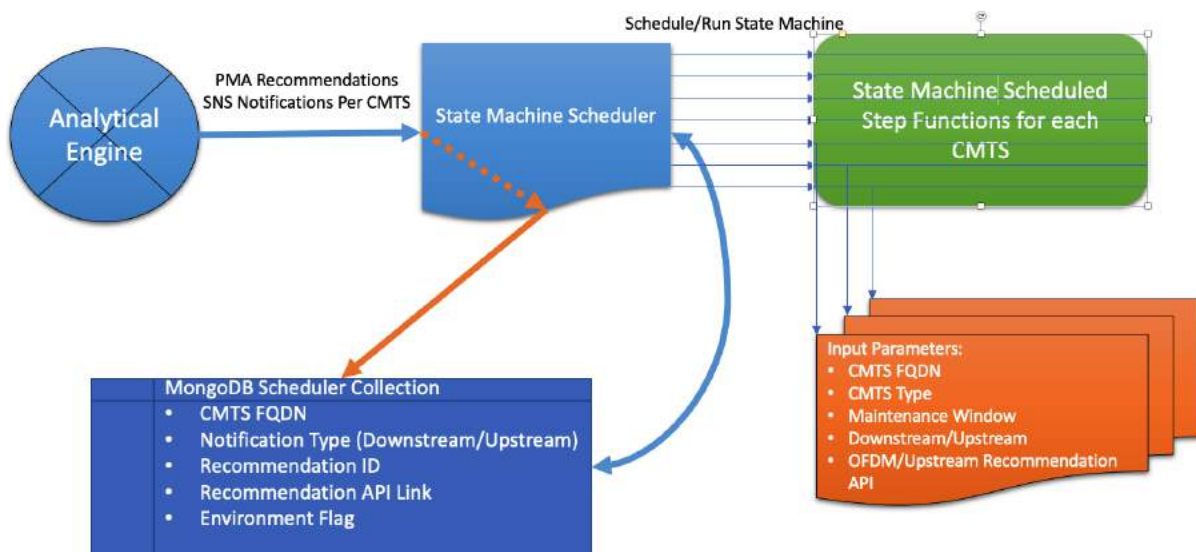
Detailed Work-flow description of Configuration Manger:



- The Configuration Manager State Machine receives a decision recommendation in a continuously timed workflow from the analytical engine. The recommendation can be to adjust or change the DS OFDM profile modulation (frequencies, QAM modulations, default modulations for each subcarrier) or can be to adjust or change the upstream profile modulation. These recommendations are based on telemetry feedback from the data engine. The recommendations should be universal in nature and are not tied to specific type of CMTS machine.
- The application schedules the state machine runs for the particular CMTS according to a predefined schedule.
- At the time of state machine operation, the CM application validates the recommended changes, conducts numerous pre-analysis checks on the CMTS using telemetry information, verifies the proper current state of the machine, and validates the overall system health.

- If the validations pass, then commands will be generated based on the AE recommendation, the type of the CMTS device, and whether it is a downstream or upstream modification.
- Before applying the commands on the CMTS machine, other validation checks are made.
- Modification to the CMTS will be made and tracked for any failures during post analysis, and the system decides if any rollbacks are required or not. The system will finally report the result of the state machine run for the CMTS, which can be tracked through the database, reporting, and dashboards.

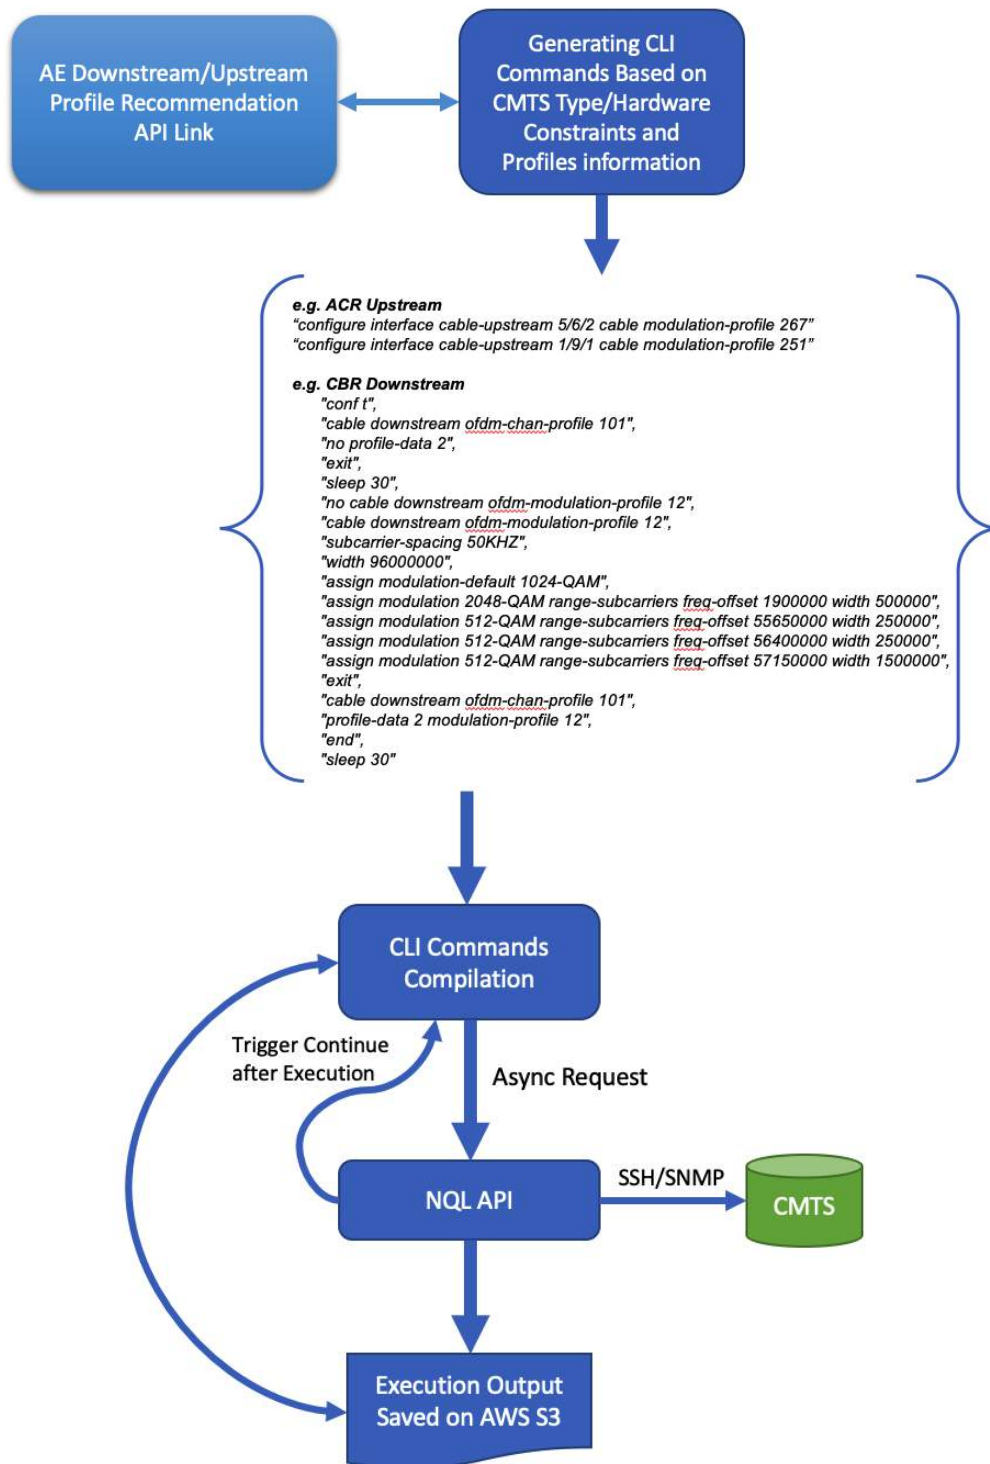
### 3.2.1 Configuration Manager State Machine Scheduler



### 3.2.2 CLI Commands Generation and CMTS Device Configuration

- In the CLI command generation phase, the function generates arrays/objects containing list of the commands which will be executed in order on the CMTS device.
- The method with which the commands are generated depends on the type of device, as each device uses its own set and way of CLI configuration. It is also based on the profile recommendation received from AE. For example, modulations and frequencies are calculated in real time based on OFDM segment start/stop, modulation bits provided, and other parameters such as lower and higher OFDM frequencies spectrum on a particular interface.
- The commands are saved in AWS S3 for the current step function to be executed during apply configuration phase in the state machine.
- In the apply configuration phase, the application uses the S3 file data object to compile the list of commands to be executed by NQL API

- An Async request is then sent to NQL API to execute the commands, the state machine waits a trigger back from NQL API for the full execution output for all CLI commands sent.
- The CLI commands output is analyzed for any errors, failures, or timeouts during execution on the CMTS.



### 3.2.3 Configuration Manager State Machine Supports and Uses the Following Technologies:

- AWS Serverless Step Function Architecture:

Each step of the state machine is handled with separate lambda function. The result of each lambda function dictates the next step of the state machine

- AWS DynamoDB  
Used for device inventory collection  
Environment flags  
Device Constraint Information
- AWS S3  
Used for dynamic command generation  
Pre/Post Analysis Device output
- MongoDB:  
State Machine Scheduled Collections and Reporting
- NetScout API for Analysis of 911 Calls in Progress on CMTS
- ServiceNow API for Change Management
- ASYNC Apply Configuration and Device Checks

## 4. Challenges / Discoveries

### 4.2 Genome

Major improvement of genome is to move from the Initial old solution with physical servers to a much more efficient and performance oriented solution.

#### Iteration # 1 Moving from old solution to Improved process – Tech stack change

Financial Implications:

- Resources
  - Had to come up with new resources who suits and understands new technology
- Infrastructure
  - Build most scalable and maintainable solution
  - Select and move to the technology stack which is cost effective
  - Secured solution
- Standardization of network
  - Support all different versions and different OS's differences
  - Solution which supports, evolving network with minimal changes
- Data reliability
  - Techniques in place to make the data reliable
  - Ways to track the data and blockers through dashboards

Political implications

- Sell the business case and architecture solution with the right reasons
- Show the brighter side of the new solution than seeing reasons how old solution didn't do
- Take the limitations had before as requirements
  - Performance improvements
  - Data Reliability
  - Dashboards
  - Alarms and notifications to act on

- SLA's in place to support downstream systems
- Long term vision of scalability and reliability
  - Minimal investment
  - Minimal efforts in upgrades
  - Much more configurable

## **Iteration # 2 Moving improved process to stable platform**

Considerations around the Stable platform:

- IPv6 protocol compatible
- Migrating from SNMP v2 to V3
- Supporting both SNMP V2 and V3
- All new solution around USM Key store

Advantages presented with the new stable platform:

- Cost effectiveness
- Reliability of network connectivity
- Security enhancements
- Enterprise level USM key store solution
- Long term vision of scalability and reliability

## **4.3 Configuration Manager**

Most of the challenges we encountered was related to CMTS hardware limitations, configuration or CLI limitations, and cable modem bugs/firmware issues which cause modems not to behave as expected with upstream or downstream profile modifications.

Iteration #1 Arris ACR E6000 Downstream:

- CLI Timeout Enhancements
- 911 Check Enhancements
- Validation Enhancements
- Working on improving overall state machine work flow for scalable deployment

Iteration # 2 Arris ACR E6000 and Cisco CBR8 Upstream:

- Cisco CBR8 implementation was challenging, during testing we found out that after modifying upstream channel profile IDs, the SNMP validations were incorrect and profile were not aligned, after significant troubleshooting, we realized that CBR8 requires upstream channel to upstream controller mapping to modify proper assigned upstream profile IDs.
- Arris ACR E6000 encountered significant issues where cable modems go into upstream partial service, which was concluded to be a bug on the E6000 which was resolved with a firmware update, and also bug in few of particular types of the Motorola modems SB21x and SB61x which we had to work on some workarounds and run debug commands to move the modems out of partial service, as well as Zoom modems required firmware upgrade which was required to prevent them from going completely offline.

### Iteration # 3 Cisco CBR8 Downstream:

- Cisco CBR8 implementation required almost a complete new design for the state machine, since CBR8 doesn't support modifying OFDM profiles globally on the cable channel, each DOCSIS 3.1 modem must be moved from a particular profile and to another one in order to modify the profile that this modem is currently using, then moving the modem back to use recommended profiles, this means significant number of CLI commands had to be generated and executed on the CMTS, one thousand modems requires 3000 CLI commands along with other profile modification commands
- We were able to achieve this by splitting the commands and apply configuration into four phases in the state machine, with each phase requiring validation check, we also used ASYNC apply configuration to limit timeouts, maximum time for each phase is about an hour to hour and half which should be fine for implementation on a fully loaded CMTS with about 3-4K DOCSIS 3.1 cable modems.

#### 1. Phase one:

Shutting down internal PMA

Locking cable modems to control profile 1 and moving service flows to profile 1 (only for profiles requested by AE)

Verifying all modems are now locked to profile 1 after waiting few minutes, if any modem(s) still stuck on other profile, CM will put the profile into ignore list and will not modify its OFDM subcarriers in phase two

#### 2. Phase two:

Applying OFDM profile changes (subcarriers and default modulations) and verifying configuration lines are applied

#### 3. Phase three:

Unlock all modems in phase one from profile 1 (moving service flows from profile 1)

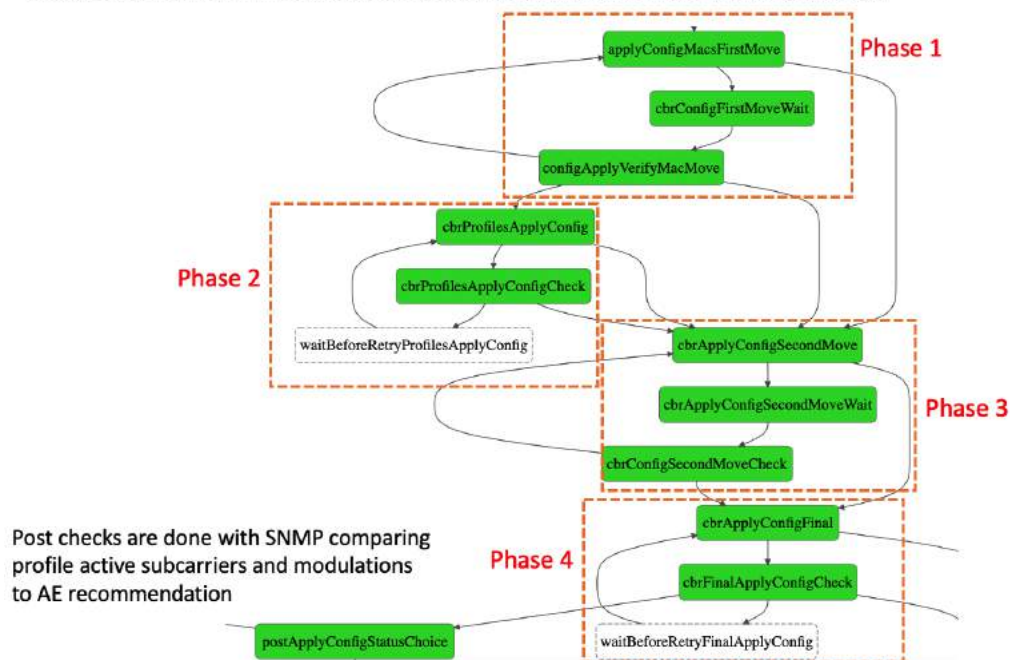
#### 4. Phase four:

Turning internal PMA back on

Do OPT commands on all modems from phase one to speed up process of modems to use appropriate profiles



Phase 3 and Phase 4 will run regardless of Phase 1 and Phase 2 result, this is basically rolling back the modems to their original state and letting the CMTS assign profiles, this must be done even if we have failures in phase 1 or phase 2, or modems will stay locked to profile 1 and will stay locked for subsequent PMA runs unless manually unlocked.



## 5. The Future

We have built a platform which is scalable and maintainable. Where do we go in the future? What are the possible improvements and upgrades that we could add to make the platform work even better for the organizations? How could we use the data that is being collected beyond the initial reason it started?

What else can we do with the configuration management platform to go beyond applying configurations? How could we improve our processes?

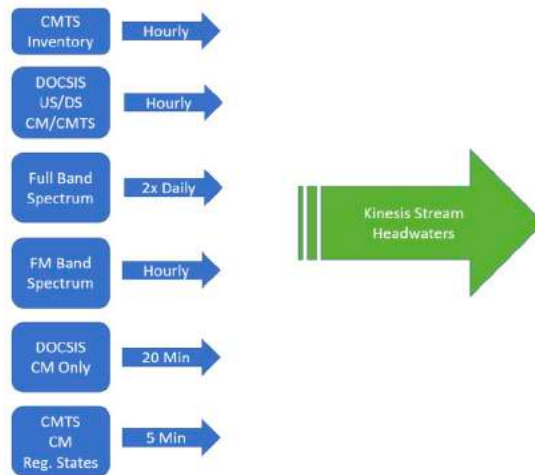
We will now share a couple of ideas and architectures for our future work in the areas of data collection and configuration management. Irrespective of the data collection side or the configuration side, major focus for future for the platform as a whole would be:

- Vendor agnostic
  - Customizable and configurable metadata for data collection points
  - Irrespective of the vendor
  - Configuration application API
  - Configuration template management
  - Upgrades to all kinds of OS upgrades and patches
- Vendor specific
  - Standardized and easy update and upgrade of vendor specific data points

- Upgrades to all kinds of OS upgrades and patches
- Easily updatable configuration templates methods

### **Data Collection: (Not Limited to)**

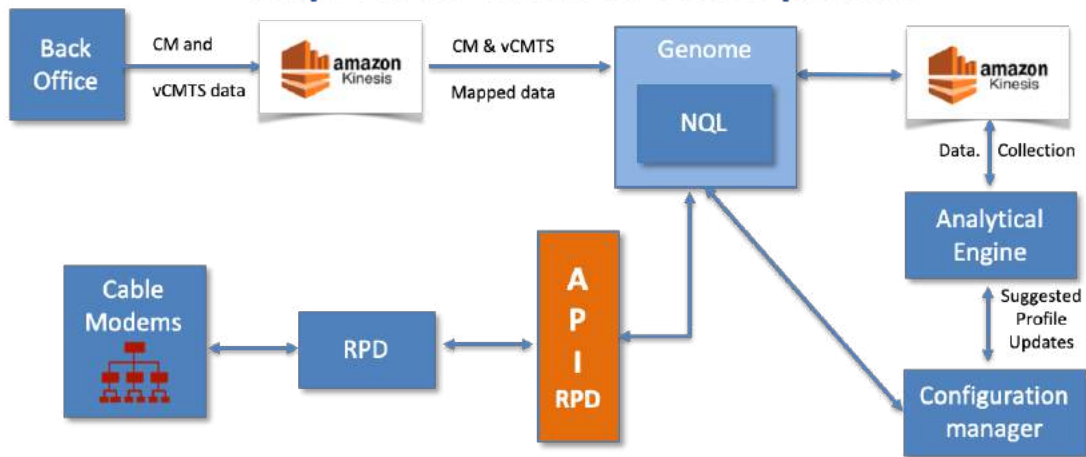
- Consolidation of the data collection across the organization
- Reduce the number of pollers connecting to devices
- Improve the performance of the devices with less pinging
- More optimizations for better performance
- Expanding the data collection beyond CMTS and CM to power supplies, RPD's ,vCMTS
- Expanding the data collection specific with adjustable frequencies



### **Configuration Management: (Not limited to)**

- Standardization of the configurations across organizations make deployments easy
- Separate it into a standalone service that so that other services can take advantage of the validations and checks it employs
- Standardize managing configuration templates
- No manual errors in initial system configurations
- Site specific customized configuration templates
- Vendor agnostic platform
- Easily adaptable, manageable, expandable

## Genome/Config Management Architecture future for vCMTS Proposed for Octave on vCMTS platform



## 6. Conclusion

Reliability, Manageability, Consistency, Usability, Performance, and Scalability are the most popular buzz words in the market, and they are now becoming the base necessary requirements for the creation of any software product/platform. Keeping that in mind, this study provides valuable insight to the ongoing evolution of creating cutting edge technology solutions for the Cable Industry.

In addition to a review of the past and present, the paper also demonstrates the benefits of using modern tools and infrastructure components. In the opinions of the authors, reliability in any kind of service/platform is the most critical component to create customer satisfaction. A commitment towards the better Customer Satisfaction always results in the best product.

## Abbreviations

CMTS	Cable Modem Termination System
CM	Cable Modem
US	Upstream
DS	Downstream
NQL	Network Query Language
ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers
USM	User-based Security Model
Auth	Authentication
SNMP	Simple Network Management Protocol
Genome	Data collection platform Name

## Bibliography & References

*A Machine Learning Pipeline for D3.1 Profile Management*, Harb, Ferreria, Rice, Santangelo, Spanbauer, 2019

# **Augmented Reality for Network Visualization**

**(Or: A Cool New Tool to Find Hidden Stuff!)**

A Technical Paper prepared for SCTE•ISBE by

**Joshua Seiden**

Vice President

Comcast Innovation Labs

4100 E. Dry Creek Road, Littleton, CO

303 486-3619

Joshua\_Seiden@comcast.com

**Nishesh Shukla**

Software Development Engineer

Comcast Innovation Labs

4100 E. Dry Creek Road, Littleton, CO

303 263-8121

Nishesh\_Shukla@comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
1. Introduction.....	3
2. A Brief History of Plant Element Visualization t .....	3
3. Realities & Challenges of Visualizing the Physical Plant with AR .....	4
4. Findings and Next Steps: .....	7
5. Conclusion.....	8
Abbreviations .....	8

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - The AR-based VON app indicating a vertical plane occlusion.....	5
Figure 2 - The AR-based VON app with occlusion turned off.....	6

# 1. Introduction

One of the higher-ranking intangibles on any field technician's wish list is the desire to spend more time fixing things, and less time finding them. In an industry with physical network assets measured in the hundreds of thousands of miles, in all imaginable terrains and environments, locating physical elements can be as challenging as their repair or replacement.

For many years, network technicians have used computer-aided design (CAD) maps to help them locate underground cables, aerial spans, taps, amplifiers, nodes, and other network elements. While helpful, CAD maps are two-dimensional, static and not always updated with current coordinates for the components listed – let alone Mother Nature's contributions, in the form of nightfall, snow-covered taps, shrubbery-occluded pedestals, or spans lost in heavy tree foliage!

Simultaneously, the landscape of Augmented Reality (AR) is emerging as a beneficial tool for just such circumstances. Not, as the reader may first connote, with special glasses, or helmets, or anything of the sort, but with a tool technicians use routinely in the field: A smart phone or a tablet. After extensive and ongoing laboratory reviews of various "smart glasses" and "AR helmets," designed to help technicians to access expert help in a "see what I can see" environment, our conclusion remains that such devices need a few more design cycles – and cost reductions – before they will represent a plausible field tool.

Instead, in 2019, Comcast Labs began to investigate the use of an AR modality as a network visualization tool that technicians can access with the same field tools (e.g. iPhones and iPads) that are part of their day-to-day lives. It applies Apple's open sourced ARKit framework to the integrated camera in an iOS device (iPhone or iPad), then draws upon graph database information, derived from latitude/longitude and GPS coordinates of plant elements, to present a near-real-time visual guide, regardless of time of day (or, especially, night) – there's a pedestal within 20 feet, but it's behind something tall, like a fence – as one of many examples. It gives technicians a sort of "x-ray vision" for the network, and moves the evolution of Proactive Network Management (PNM) from 2D to 3D. The tool, which we call "VON" for "Visualize Our Network," went into the field in early 2020, as a component within the tool suite technicians use daily.

This paper will describe the VON work to date, including the evolution of network visualization, how AR works to visualize the plant, development challenges, why graph theory was an important design component, the importance of "occlusion," lessons learned, and next steps (if only in the form of a "design wish list.") Note: "Occlusion," which will be mentioned frequently in this paper, is defined here as "the blocking of a view of part or all of something, but something else."

## 2. A Brief History of Plant Element Visualization

Network visualization, in a classic two-dimensional, addresses-on-paper and rolled-up-maps sense, has long been a vitally important, if passive, component of proactive network maintenance (PNM). Many "cable old-timers" recount stories of entering the industry via a summer job – "here's a clickwheel, and here's how to use it" – to measure, verify and/or document the physical topologies of the network.

The state-of-the-state is now digitized, and revolves around spatial mapping in graph form. A Graph Database is a database that not only stores the elements of data, but also stores the relationships between those elements. The data is actually stored in the graph in the same way that they are connected in the "real world." Graph, in this sense, means "visibly connected," in the sense of showing the relationships between mapped elements: The amplifier goes to the line extender, which connects to the tap, then the splitter, and so on. Graph databases are the latest in what's been a long progression, happening in parallel

with other PNM initiatives, and all in the pursuit of identifying and resolving plant issues before they impact customers. Graph databases will be discussed further in Section 3.

Most (but not all!) operators long ago mapped their network topologies into formats that can be reasonably easily translated into graph models. It was a natural evolution from computer-aided design (CAD), back when the state-of-the-art for network visualization were lists of addresses, which could be assembled into problem clusters.

Geographic Information Systems (GISs) came next, which yielded “dots on a map” informed by Global Positioning Data (GPS) coordinates. GIS/GPS advanced network visualization into tree-and-branch representations, which showed subscribing households connected to the branches. Since then, Google Maps emerged as a powerful visualization tool, but even in “satellite view” it’s largely ineffective in seeing anything smaller than, say, a car in the driveway. Certainly not a tap; maybe a pedestal, depending on seasonality and camera angle at the time the image was collected.

The next step in network visualization, fueled by AR, takes matters one important, three-dimensional step further. Previously, the “dots on a map” showed the connection between, say, an amplifier and a tap, using each element’s GPS-based latitude (X-axis) and longitude (Y-axis.) AR brings a new dimensions – a Z-axis – which adds visual depth and the camera attitude (pitch, yaw and roll) for first-person viewing perspective. This added depth allows the technician to do something that they’ve not previously had the capability of doing – visualizing multiple elements of the network, overlaid on the “real world” simultaneously. Not quite Superman-grade x-ray vision, but close enough to draw the parallel!

### **3. Realities & Challenges of Visualizing the Physical Plant with AR**

As a painfully basic observation, Augmented Reality, as a technology modality, augments actual reality with informational overlays. And, as it turns out, overlaying the “this” onto the “that” – meaning most types of ancillary information designed to augment a reality – isn’t the hard part of AR. The core technologies, like ARKit, are largely open-sourced, and substantial amounts of documentation exists online. This makes the actual “overlay” part of AR implementations reasonably straightforward.

The first reality when building an AR model of the network is the need to build a graph database of the network. If each element in the graph has latitude and longitude coordinates, and we understand the relationships of each element to one another, we can represent the network in 3D space. This brings us, again, to a discussion about the necessity of graph databases when developing AR-based plant visualization techniques for technicians.

Traditional relational databases rely on tables, where the data are organized into well-defined columns and data formats. This system works well when dealing with a modest number of relationships between a few different tables. But traditional relational databases are woefully underpowered to meet the needs of visualizing all of the components present in a modern broadband network.

A graph database is one that organizes data according to the elements of a graph – with its own “nodes” and “edges” that represent the relationships between the edges. Graph databases represent network connections in real time, which closely resemble how we interpret them naturally. As such they are ideal for informing network topologies and related maps in a visual way.


That said, there is the matter of the “Z-axis,” mentioned in the previous section – another reality encountered when developing an AR-based plant visualization tool. Creating an AR session necessarily includes X, Y and Z coordinates: Latitude, longitude, and in this case, depth. Depth and attitude augmentation were a challenge, in part because ARKit was in beta when we began this effort, and partly



because our use case required continuous depth detection, while in motion, to direct a tech to the desired destination. We added the Z-axis by using an ARKit feature called *plane detection*, and in the case of the Visualize Our Network AR effort, vertical plane detection.

Consider the technician looking for a pedestal in an established neighborhood, which is to say “in full foliage,” because it’s summertime. Or, that same neighborhood, buried under a foot of snow in February. Either way, the pedestal in question is behind a fence. To handle those frequent cases where a network element of interest is occluded by something – a fence, a tree, a building – we detect and notate the occlusion by marking it blue. This tells the tech that the object of interest is behind something vertical. We determine occlusion by detecting vertical planes in front of the object of interest.

Without vertical plane detection as a means of marking occlusions, it would be very hard, if not impossible, to pinpoint a specific location, in 3D. Figure 1 shows the VON app with Occlusion turned on, to indicate a vertical plane obstruction. Hence, any objects beyond the Occlusion would not be visible, such as pedestals, cable, and other elements of our plant.



**Figure 1 - The AR-based VON app indicating a vertical plane occlusion**

Figure 2 (next page) shows the same vertical occlusion – the same fence – with occlusion turned off. This allows the technician to view the elements of our infrastructure, including the vertical occluded objects, but without noting them as such. This combination gives a visual depth to VON that previously didn’t exist.



**Figure 2 - The AR-based VON app with occlusion turned off**

For underground plant, we applied a “snap-to-ground” focal point for the VON app, which, in essence, extended a “negative vertical plane” of -6 feet. As a result, by starting at a node, then going “down” the graph database to the first tap, then the second, and so on, we were able to make the assumption that if Node A is connected all the way down to those seven taps, there’s likely a buried cable between them, which at best is observable as not being aerial, but rather, underground, indicated by a line/plane on the ground. Without a graph database, that job would’ve required a physical walk-out.

So: The hard part of AR-based plant visualization, in practice, isn’t the AR. The hard part, perhaps not surprisingly, is assembling the information correctly: Latitude, longitude, depth, what kind of object, and whether it’s occluded by another object. An early and persistent challenge was simple GPS inaccuracies, either from a faulty phone/app interpretation, or from incorrect GIS entries. Truth be told, we encountered at least one pedestal that was, according to its GPS coordinates, in the middle of a lake!

Another challenge: Point of view. While AR technologies continue to advance, there’s still a lot of room for improvement. For instance, a marker representing a pedestal on the phone’s or iPad’s screen may appear to be three feet away on the screen – but actually, it’s 30 feet away. The reason is the perspective: The phone is stitching a camera feed together with a virtual object. The virtual object has no explicit knowledge of what the camera is seeing. Rather, it exists in an empty world, with only an X (latitude), Y (longitude) and Z (depth) axis as a landing spot. If the perspective is off, while you’re walking down an alley, looking for a plant element, it can *look* like you’re close, when you’re not. Resolving this took some iterations (understatement)! While tying a virtual object to a scene could be achieved with the use of model detection and logic code, to detect and identify scenes, such a plan would rapidly deplete the phone or iPad’s battery.

A third challenge involved the previously mentioned matter of occlusions encountered along the way, and especially vertical occlusions – the fence, with the pedestal behind it. The VON, without scene detection of some kind of understanding of the surroundings, can't know that there's a "fence" in the way, only that there's a vertical plane occlusion. We resolved this by incorporating a way to adapt the vertical plane, to show / not show a vertical object in the way of the desired plant element. This discovery went a long way in making the user interface look realistic, and created a means for the technician to interpret depth. Without the Occlusion information, it doesn't make sense, because it lacks visual depth.

## 4. Findings and Next Steps:

The core technology enablers that made it possible to develop the VON application, as a PNM tool, were ARKit 3 Beta (at the time), GIS, and Graph Theory/Database. It is important to again note that the graph database aspect was crucial to VON as it described the relationship between each element of our infrastructure and how it all comes together. Graph theory, in combination with the GIS data that we had on our infrastructure, allowed for a graph database to exist and then be used by VON to visually represent all the different elements.

ARKit 3 was in a Beta state from Apple at the time we embarked on the VON tool. Because of the newness of the release, it allowed VON to tap into the potential for occlusion, and how it can beneficially alter the perspective of a user. Specifically, in terms of marking and presenting the visual depth of the geographical virtual objects, portrayed ontop of a digital camera feed. This is also known as "Markerless Augmented Reality," defined as AR that doesn't need prior knowledge of a user's environment in order to overlay 3D content into a scene and tie it to a fixed point in space. After the introduction of ARKit 3 in mid-2019, creating AR Sessions with virtual objects based on geographical data acquired from graph databases was straightforward. The challenge (still) is the available plant data and the GPS inaccuracies of sensors within our handheld devices.

Simply put, accuracy matters to visualizing the physical plant with AR. It matters in particular as to how we place the virtual element, onto the digital (camera image) feed. Both the stored geographical data about our infrastructure elements, and the GPS circuitry present in handheld devices, need a reasonable degree of accuracy to best inform the VON effort. Currently, inaccuracies can create a visual drift from the intended object of three to 12 feet. This is why a favored new feature for VON is the ability to re-tag a network element with correct location coordinates, quickly and easily, and as technicians use the app to locate objects. Better data = better accuracy = finding hidden elements faster.

To more thoroughly test the VON application, in real-world circumstances (no pun intended), it was "easter egged" into an existing suite of PNM apps used by Comcast's field technicians. In that sense, it's "out now," although truth be told, only a slight percentage of the field has yet to find it and put it to work. Also, it requires an iOS device that is optimized for ARKit 3 and above (iPad mini, iPhone 7 and above), which not all technicians have, as a function of routine device upgrades.

As with any new Augmented Reality application, the initial reactions run the gamut from "how cool" the application is, to an overall and genuine excitement about it. The excitement about it will drive more technicians to use the application, and provide feedback to improve it. As we gather this feedback, and continue to experiment with the "art of the possible," a few items tend to rise quickly to the list of desired feature additions for VON.

Most involve giving technicians more information about a task, by providing historical context. For instance, the ability to "hover" over a home, then "see" as an AR overlay its known technological history – CPE spectrum measurements, or full spectrum analyses from gateways, as two of several potential

examples. From a plant element perspective, the ability to augment, say, an amplifier, with a visual overlay of its latest sweep trace measurement, would be a time-saving and useful troubleshooting tool for the field.

Another and more practical/tactical consideration is an easier, on-the-spot/in-the-moment way to correct incorrect database entries, as mentioned above, by visually “attaching” to a plant element, and thereby correcting bad data.

## 5. Conclusion

This paper described an ongoing effort within Comcast Labs, and active in the field with technicians, to use Augmented Reality for finding specific plant elements – and especially those that aren’t plainly obvious, either because it’s dark out, or because they’re covered up by something else. Called “Visualize Our Network,” or VON, it uses AR to add a 3D element to Proactive Network Management. It very specifically uses an iPad for the visual element, and not “AR glasses,” which can get lost, or be unwieldy for people who already wear corrective eyewear.

Equipping technicians with this kind of 3D “find it” tool is both approachable and very well received, both for the on-the-job assistance, and the “cool factor.” Using open-source elements like ARKit, and assuming an accurate plant database exists, with latitude and longitude coordinates for key plant elements, like nodes, amplifiers, and taps, you too can trick out your PNM tools with 3D plant visualization!

## Abbreviations

AR	Augmented reality
ARKit	Augmented Reality Kit
CAD	Computer Aided Design
GIS	Geographic Information Systems
GPS	Global Positioning Satellite
PNM	Proactive Network Management
VON	Visualize Our Network

# Verification of Electrical Grounds/Bonds Using Computer Vision

A Technical Paper prepared for SCTE•ISBE by

**Shawn Kercher**

Principal Engineer  
Comcast Innovation Labs  
4100 E. Dry Creek Road, Littleton, CO  
Shawn\_Kercher@comcast.com

**Jacob Hallberg**

Software Development Engineer  
Comcast Innovation Labs  
4100 E. Dry Creek Road, Littleton, CO  
Jacob\_Hallberg@comcast.com

# Table of Contents

Title	Page Number
1. Introduction & History .....	3
1.1. A Brief History of Electrical Grounding and Bonding .....	3
2. Job #1: Images & Training Data.....	4
3. The Other Job #1: Annotating Images .....	7
4. The App & How It Works .....	9
4.1. WebApp Version .....	9
4.2. iOS App .....	10
4.3. Trial .....	12
4.4. Future updates .....	12
5. Conclusion.....	12
Abbreviations .....	13
Bibliography & References.....	13

## List of Figures

Title	Page Number
Figure 1 - A selection of images from the House Box category.....	5
Figure 2 - Image example of House Box category; panels 2 and 3 highlight the bond block and filter, respectively .....	5
Figure 3 - A selection of images from the Power Bond category.....	6
Figure 4 - Each panel shows the bonding hardware highlighted .....	6
Figure 5 - Example annotation of a House Box filter object of interest.....	7
Figure 6 - An active learning pipeline to reduce costs (time and money) of image annotations .....	8
Figure 7 - A view of the WebApp version for Computer Vision grounding/bonding.....	10
Figure 8 - A view of the iOS app for Computer Vision grounding/bonding .....	11

# 1. Introduction & History

Electrical grounding and bonding, as a means to provide a low-impedance electrical connection between two or more metallic bodies that are normally not current-carrying, has long been a necessary task for the industry. Electrical grounding prevents circuit overloads and removes dangerous ground-fault voltage on conductive parts. Primarily, grounding and bonding is a mechanism to protect homes and businesses, and their contents (including people), from electricity surges, such as from lightning.

When all of the communication and power grounds are working at a home or business, and are all “good grounds,” any stray current in the system goes to ground and no voltage differences exist between them which is the desired state. However, when grounds are bad, the current can and will find its way into the CATV network. When voltage appears in the CATV network the voltage potential problem appears to “drop the voltage” of anything connected to the network. Set-top boxes, modems, gateways, or, in the case of large MDUs, line extenders all have seen the drop in voltage. Stories abound from technicians taking the hit of 90-volt shocks from power-passing taps, or 120-volt shocks when hooking cable up to taps in MDU lockboxes – which makes for a decidedly unpleasant site visit! Proper electrical grounding and bonding is as equally or more important to business/enterprise customers especially when lightning hits a facility, knocks out half of its servers, and turns out to be a grounding issue – as in, “our fault” (pun intended.)

This paper describes the use of Computer Vision (CV), Artificial Intelligence (AI) and Machine Learning (ML) to provide technicians with a visual verification that gives them a high level of confidence that the structure (the home or business) they’re visiting is properly grounded and bonded, using an app developed at Comcast Labs, on their work-issued iPhone or iPad. The paper covers the background and history of grounding and bonding; why those activities continue to be vitally important in communications network performance; and the importance of images and image annotation to build a production-grade database of proper grounds/bonds that can be put to work by field technicians when assessing a home or business.

## 1.1. A Brief History of Electrical Grounding and Bonding

While electrical grounds and bonds have been overseen since the late 1960s by entities like the National Electrical Code (NEC), in part 820 [1], compliance reporting is generally not a requirement. Unlike rules like Cumulative Leakage Index (CLI), which require quarterly compliance reports, the job of making sure the ground/bond is true is voluntary and not enforced. That said, some cities and franchise authorities have instituted grounding and bonding inspections, to monitor how many jobs are properly grounded and bonded. Improper grounding/bonding in those instances have incurred fines, which is to say it also impacted profits.

Our work to automate the inspection process for proper grounds and bonds began as an offshoot of a related project, which uses CV, ML and AI to catalog the contents of racked equipment in data centers. The intent was to automate regular equipment inventories and establish a visual map of where servers are, physically, in what can be large and cavernous facilities full of such

gear. Late in 2019, we were approached with the idea of extending the effort towards outside plant activities, in particular, helping technicians to verify that the electrical ground is good. Considering the fact that electrical grounding activities began before WWII, and that the U.S. counts roughly 125 million homes (a number that does not include businesses), the range and types of electrical bonds we encounter are vast. Countless (and uncounted) variations exist in the field. Hence the importance of training data and the establishment of a solid visual database to catalog existing manners of grounds and bonds.

## **2. Job #1: Images & Training Data**

The reason why we're talking about good electrical grounds and bonds in the year 2020 is the as network operators we are responsible for the safety of our technicians, customers, their homes/businesses, and the electrical devices within those homes and businesses. Merely through expectations and the accountability that we take each day, in connecting customers to our network, the electrical ground and bond is a vital, if sometimes overlooked, activity. After spending the last decade stabilizing the major portions of the network – platforms, headends, CMTS devices, outside plant components – another domino in the lineup, in the pursuit of overall network reliability, is grounding and bonding.

This effort started by developing an app, accessible on an iPhone, iPad or web browser, with linkages to images of about 30 grounding/bonding implementations encountered in the field. The app was given to roughly 50 employees, at the start, to capture and ultimately catalog a visual database. These employees were spread out over the United States, giving us a varied dataset of 1,500 bonding and grounding images. This varied dataset helps combat the identification problem of the many different types of grounds and bounds that stem from WWII-era home construction onwards.

In order for the system to be production-grade from a Computer Vision perspective, substantially more training images will be required – as many as 10,000 or more. This estimate can be drastically increased if proper selection of images is not taken. For example, the bonding hardware, the bond block, and filters can all vary in appearance depending on the region of the country and when the home was built. This data variance makes it harder for the system to learn all of the different types of bonding and grounding hardware it will see in the field. Consequently, if not addressed correctly, the system will require additional training images to be production-grade.

The visual database is split into two categories: house box and power bond. The house box category generally consists of a frontal shot of a house box. The house box is in most cases mounted to the side of the home and houses the bond block and the Multimedia over Coax Alliance (MoCA) filter (see Figure 2). The power bond category generally consists of many differently angled shots that capture the bonding hardware (see Figure 4). Categorizing the images in the visual database into these two categories is an essential part of training a production-grade system. It allows the system to learn the specificities of each category, improving accuracy and robustness.



Currently, our bonding and ground AI uses approximately 1,000 house box images and 500 power bond images to train the system. Using machine learning techniques called Transfer Learning, Data Augmentation, Regularization, and Semi-automated Image Selection, we are able to drastically reduce the total number of images needed for training. Based on our current rates of accuracy for both power bond and house box, we estimate we'll need upwards of 10,000, or more training images to release a production-ready system.



**Figure 1 - A selection of images from the House Box category**



**Figure 2 - Image example of House Box category; panels 2 and 3 highlight the bond block and filter, respectively**




**Figure 3 - A selection of images from the Power Bond category**



**Figure 4 - Each panel shows the bonding hardware highlighted**

### 3. The Other Job #1: Annotating Images


Once the visual database is established, and growing, a corollary and vital task is image annotation – to manually mask the objects of interest within an image. As important as the image itself, image annotations are the backbone of training a Computer Vision system. Annotation requires a person to open the image in an image annotation tool and “mask over” the object of interest with a polygon shape (see Figure 5). All pixels within the polygon shape are then labeled to correspond to the object of interest. In bonding and grounding, the objects of interest are the bonding hardware, bond blocks, and filters. Note the points that make up the polygon mask. All pixels within this mask are deemed to correspond to the filter object for this specific image, which is used to train the Computer Vision system.



**Figure 5 - Example annotation of a House Box filter object of interest**

This annotation job must be done for each and every image that enters the visual database. Additionally, each annotation must be checked for accuracy, to ensure that any bad annotations do not negatively affect the accuracy of the system. After the images in the visual database have been annotated, and their annotations verified, they are finally ready to be shown to the Computer Vision system to learn the intricacies of the object.

To ease the burden of annotation, we made use of active learning (see Figure 6). Active learning is the process of training our Computer Vision system, first with a small subset of labeled images, and then using the system’s learned ability to decide which new subset of images in the visual database should be annotated manually. We then annotated the new subset of images from the system manually, and repeated the process. This process of human-machine iteration is essential to reducing the cost of annotation, and overall improves the end result of system.



**Figure 6 - An active learning pipeline to reduce costs (time and money) of image annotations**



## 4. The App & How It Works

To approach the job of creating a Computer Vision-enabled means of visually verifying electrical grounds and bonds, we built two apps: a WebApp and an iOS App. Both use the same API endpoint to upload images to be inferred to the Computer Vision model. This process leverages GPUs for detecting Filters and Ground Bonds. The CV model is coupled on what it detects, using rules to determine if a filter and ground block is wired correctly. Each image is captured, saved to our database and added to the ML training data. A feedback section indicates whether or not the CV model detected the image correctly.

Rules were developed to determine a simple validation for the house box and the power bond, as characterized below:

### House Box

- Does a filter exist?
- Does a bond block exist?
- Bond block connected to filter?

### Power Bond

- Does a tag exist?
- Does bonding hardware exist?
- No prohibited bonding locations?

### Steps:

1. Snap a picture or upload an image of the House Box or Power Bond
2. Receive an answer if Ground Bond/Filter is properly installed

### 4.1. WebApp Version

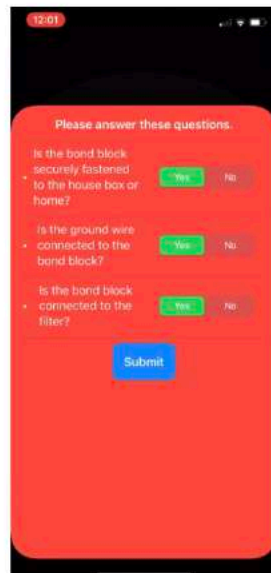
This version was built specifically to run on a phone browser. We started with a WebApp to iterate quickly. This tool can be exceptionally valuable for testing (see Figure 7.)



## House Box




1 Take Picture



2 Answer a few questions about the Bond and Ground



3 ML CV model identifies the Filter and Ground



4 Validation

## Power Bond

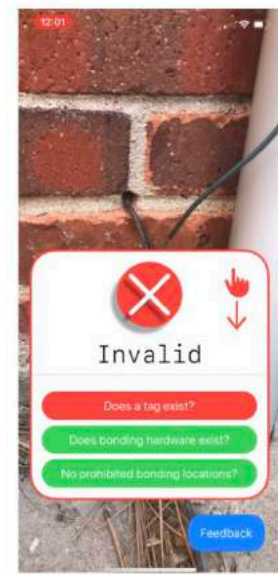
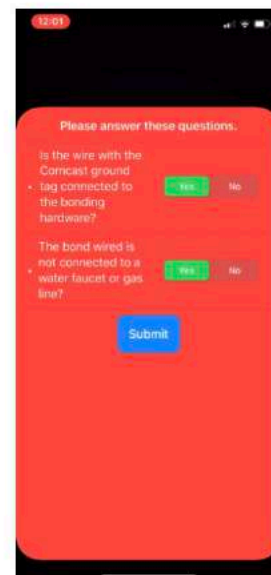



Figure 8 - A view of the iOS app for Computer Vision grounding/bonding

### 4.3. Trial

The Computer Vision app for electrical grounding and bonding verification is currently in trial in all three Divisions of Comcast, and specifically in systems within our Georgia, Florida and Colorado footprint. The app was successfully deployed to 50+ technician iPhones, used by technicians who are participating in the trial. So far, our feedback and learnings have been crucial to gain knowledge on how technicians are using the app, and to interpret different use cases we have not seen before.

Next steps are to integrate the app with an internally-developed digital tool, developed by and for technicians, called “Tech 360,” and designed to give them a 360-degree diagnostic view of the homes and businesses they visit. That work is slated to be complete by year-end 2020, with anticipated availability to our 35,000+ techs by the middle of 2021.

### 4.4. Future updates

As the Computer Vision app expands in reach and usage, we are compiling a list of desired feature additions for the WebApp and iOS app. The list below represents a partial summary:

- *Development of an Image Queue:* If no mobile or WiFi network is available, enable the app to cache the images, then retry when connectivity resumes
- *Ease connection procedure:* Implement SSO and remove VPN
- *Implement wire tracing:* This will evaluate both RF and ground wires to ensure they are properly attached
- *Implement logic:* To determine if bond block is attached to the house box
- *Increase usability:* Add flags and notifications
- *Increase distinguishability:* Implement the ability to distinguish between six different filters that are deployed in the field.

## 5. Conclusion

Assuring the existence of a proper electrical ground and bond is a vital and longstanding best practice in cable telecommunications and broadband networks. What makes the task challenging is the simple fact that electrical grounds and bonds date back to the earliest days of homes with electricity. They’re necessary to provide low-impedance electrical connections between two or more metallic bodies that are normally not current-carrying, to prevent circuit overloads and remove dangerous ground-fault voltage on conductive parts.

Good grounds mean that any current in the system goes to ground, and no voltage differences exist between them. Bad grounds can create voltage potential problems that appear to “drop the voltage” of anything connected, whether set-top boxes, modems, gateways, or, in the case of large MDUs, line extenders. Bad grounds can also wreak havoc when lightning strikes.



This paper described the use of Computer Vision, Artificial Intelligence and Machine Learning to provide technicians with a visual confirmation, intended to give them a high level of confidence about the ground and bond of a home or business. Using a WebApp and iOS app developed at Comcast Labs' Denver-area facility, technicians can verify grounds and bonds on their work-issued iPhone or iPad. It covers the background and history of grounding and bonding, why those activities continue to be important in communications network performance, and a discussion of the image collection and annotation necessary to build a production-grade visual database of proper grounds and bonds. This effort is in trial now (summer 2020), with plans to integrate this technology into a core digital diagnostics app used by field technicians following a successful trial.

## Abbreviations

AI	Artificial Intelligence
CATV	Cable Television
CLI	Cumulative Leakage Index
CMTS	Cable Modem Termination System
CV	Computer Vision
GPU	General Processing Unit
MoCA	Multimedia Over Cable Alliance
MDU	Multiple Dwelling Unit
ML	Machine Learning
NEC	National Electrical Code
RF	Radio Frequency
SSO	Signal Sign On
VPN	Virtual Private Network

## Bibliography & References

[1] National Electrical Code, Article 820: Community Antenna and Radio Distribution Systems: <https://www.ecmweb.com/national-electrical-code/code-basics/article/20885527/article-820-community-antenna-tv-and-radio-distribution-systems>

# Aggregate Wi-Fi Telemetry Use Cases

A Technical Paper prepared for SCTE•ISBE by

**Colleen Szymanik**  
Sr. Principal Engineer  
Comcast  
1800 Arch Street Philadelphia, PA 19103  
215-970-1953  
Colleen\_Szymanik@comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
1. Introduction.....	3
2. Data Sources.....	3
3. Metrics.....	3
3.1. Interface Statistics.....	4
3.2. Radio Statistics.....	4
3.3. Client Station Statistics .....	6
4. Collection Interval.....	7
5. Active vs. Passive Measurements .....	8
6. Conclusion.....	8
Abbreviations .....	9
Bibliography & References.....	9

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - 2.4GHz and 5GHz available channels .....	5
Figure 2 - Interference.....	6
Figure 3 - Client capabilities.....	7
Figure 4 - Aggregate Trending.....	8

# 1. Introduction

Assessing WiFi performance can be challenging, especially at scale. Most operators and service providers need provisions to understand the quality of delivered service to ensure customer satisfaction. In order to assess the service, we need to understand the requirements for consistent and reliable service offering over WiFi. Since WiFi can be used over both 2.4 GHz and 5GHz frequencies, they should be measured and assessed separately given there are different challenges per radio frequency. In this paper, each router that services a customer home will assume 2 connected radio interfaces. Once the 6GHz band is widespread, this will be an added (3<sup>rd</sup>) interface to quantify. Each connected client device, such as laptop or smart phone, will also have different bandwidth requirements to obtain the intended service, regardless of the connected interface. In this paper, we will talk through data sources and collected metrics. From those metrics, is it possible to make assessments at a aggregate level to infer reliable and consistent quality of service from poor or unusable service? Conversation pairings will refer to the transmit and receive between an access point (AP) and client device, such as a laptop. The idea here is to refer to this as a simple conversation. In order to have effective communication between both the AP and the station, the rules put in place by the protocols and the WiFi specification need compliance. Assuming compliance is met, measuring quality is the next step. WiFi quality depends on the specific conversation pairing requirements. Each pairing does not have the same time and bandwidth requirements, such as an IoT device compared to a 4K streaming device. They should not be treated the same when appraising the connection quality.

## 2. Data Sources

The Access Point (AP) to client station connection is typically a one to many relationship. As an operator, measuring data from the AP is typical since it is the central point of service in the home. It is also usually managed at the service provider level. The other main element in the equation is the client station device. While client station is not under the service provider control, it is important to recognize that the important role this device component has in this exercise. These measurements are solely to provide consistent and reliable connections to these client devices. While client devices have varying capabilities, it is fundamental to consider what is needed for each of the individual client types. Some of the most successful ways to rate the level of service is to be as individual as possible for all the requirements. Clients on a network have specific needs. As a service provider or operator, can we measure those needs at scale?

## 3. Metrics


Most measurements for WiFi quality of experience combine the same basic elements. The primary components consist of interface statistics, radio metrics and client connection statistics. The challenge with these elements are the fact they are constantly changing. WiFi is a very resilient protocol, as with most networking elements. These elements are constantly changing to meet the current conditions for service delivery. An example of the changing conditions is the rate adaptation mechanisms in video delivery to drop from HD service to SD service if HD service levels are not met. Adaptability is a key method for resiliency, but it poses challenges to understand when the quality of service compromises the customer experience. The components can be measured and assessed to provide feedback on that customer experience.

### **3.1. Interface Statistics**

Interface statistics are useful in understanding what was successfully transmitted and/or received between the AP and the client station. How often were those transmission successful, filled with errors, or needed to be repeated? When we are quantifying applications, such as video, some retransmissions can be tolerated due to provisions such as buffering. While video is mentioned, each application for each client device type can be tailored to the the corresponding applications. The result is tailored to the specific service level requirement for a good user experience. Having a good ratio for interface statistics can help properly assess whether or not the current level of interface statistics are viable for a consistent user experience. Interface statistics consist of sent and received, as well as errors, retries and retransmissions. One of the key differences could be fine tuning ratios that can prove service level offerings. If it is known what is transmitted and received. Looking at the proportion of frames sent and received with respect to retransmissions and errors can yield insight into the quality of the connection. As such, measuring round trip time to assess latency is another potential method to assess the service quality between the AP and the client station. In gaming, most gamers will use latency as a measure of whether or not they are able to achieve consistent service.

### **3.2. Radio Statistics**

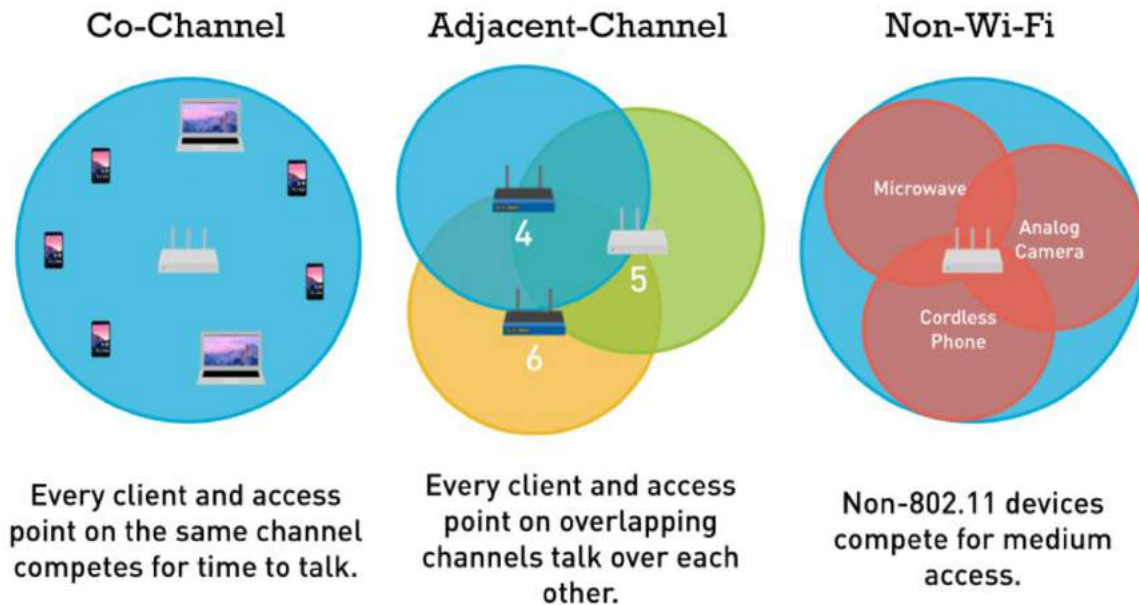
One can argue that gaining access to the medium is the biggest challenge with WiFi today. In fact, from 802.11's beginnings, access control still presents a challenge. WiFi has always been a "listen before you talk" service. MU-MIMO was the long awaited feature in WiFi5. Most of the excitement for WiFi6 is OFDMA. Efficiency in admission control to the medium is sought after in WiFi implementations. If the AP can successfully group devices to transmit and receive at the same time within the group, this can free up more airtime to be used for even more transmissions. The number of concurrently connected devices in a home are growing as more and more devices are WiFi capable. Practices such as band steering, pushing a client device on the 5GHz band is one way to make better use of a less congested frequency. The unlicensed 2.4GHz band is already over used and over crowded, especially in densely populated areas. Measuring the radio quality is key to understanding how much time is effectively being used for client connections. In 2.4GHz, there are only 3 non overlapping channels to use.



**Figure 1 - 2.4GHz and 5GHz available channels**

Given that WiFi operates in unlicensed spectrum, there is a lot of congestion in this space. As mentioned, it is typically worse in the 2.4GHz band, but congestion can still be an issue in 5GHz, especially in dense urban environments. This usually results in a lot of extra interference in this band that makes the signal to noise ratio for those connected clients lower than that in the 5GHz band. The noise from WiFi will result in interference.

## The three main causes of WiFi interference



**Figure 2 - Interference**


In both bands, measuring the channel utilization where we can understand the amount of time used for servicing connected client stations is important for overall radio measurements. While there are some ways to do this in a standardized method, there are slight nuances that make normalization between APs necessary. Depending on the device, operators may see measurements such as transmit opportunity (TXOP) or free airtime (FAT) that begins to characterize the radio interface. These differences in the measurement will require normalization in a multi vendor environment. If we treat the radio measurement in the same fashion as the conversation pairing, it is a natural progression to normalize the measurement for aggregate trending.

### 3.3. Client Station Statistics

Client connection statistics that reflect signal levels of the current connection are always present in most quality of experience assessments. RSSI, PHY rates, MCS values, SNR values all provide feedback to estimate the AP to client station connection. Once again, there are nuances for these measurements. It is important to know what exactly is being reported. In WiFi, the management frames are transmitted and received at the lowest mandatory rates. If the reporting includes these management frames, this can skew the estimate of the connection quality. It is also worth noting that each client type is limited from the physical capabilities.

## BREAKING OUT BY DEVICE (IOS)

Gig Speed Test on Different iOS Devices



- Older device are capable of less speed
- Really old devices such as iPad 2 are not capable of speed above 25 mbps
- The very best tests on the latest iOS devices reach ~650 mbps
- No iOS device has achieved full gig speed
- Tier is not a limiting factor in the tests shown

Figure 3 - Client capabilities

Not all client stations are created equal and should not be measured as such. As mentioned in Data Sources, operators can use the APs to take all these measurements so they know what is being measured and how often. Typically, APs have more transmit and receive antennas than the connected client devices. That imbalance of radio power needs to be recognized when evaluating these signal levels. Thresholds for these client types need testing feedback to ensure the levels are reflective of user experience.

## 4. Collection Interval

We mentioned earlier that WiFi is adaptive and constantly changing for current conditions. Collecting data at scale from moving targets should be considered when choosing collection intervals. Even if data storage and infrastructure is scaled well to collect these metrics often, it might not be necessary to collect data every minute or even every second to get reasonable measures of quality. In this conference, the idea of limitless possibilities is challenged. If that theme is extrapolated to the point where data storage and processing power is not an issue, do we want to collect data frequently? Does that make sense? In all computing systems, once mechanisms are built, efficiency is the next logical step, even if scale isn't an issue. Collecting data at the edge and aggregating into a cloud element is only meaningful if the right building blocks are in place. Beginning with the conversation pairing is fundamental to then build groups to count and ultimately trend and aggregate. Using that order of operations as a rule set can yield consistent aggregation points. If one can target the right thresholds for the individual conversation requirements, how often is necessary to check that connection? A good start would be hourly, if the infrastructure can support the individual pairing data model. A more granular set of metrics has been found to roll up 15 minute intervals. There are some methodologies that will have localized collection agents at the WiFi interface to aggregate the data to be collected less often by centralized infrastructure.




Strategies for this type of data is important make sense of this at scale. If one keeps averaging or aggregating data, it will just make the measurements less meaningful and will not capture the experience properly. If we can divide the data into different compartments for each measurement, we can count those data occurrences while quantifying scale at that individual level. As a result, collection intervals can be done less often and aggregation can be scaled while still providing a meaningful experience gauge. The idea here is to expand vertically rather than horizontally. Limitless possibilities doesn't always mean expanding in the same direction.

## 5. Active vs. Passive Measurements

All of the data metrics and measurements discussed at this point are assuming that these measurements are passive. This means that we are going to watch active client station connections to a given AP and measure the primary components and report those measurements accordingly. Given that applications and client stations change hardware and software often, there is a need to pulse ever changing conditions. Having active measurements can aid in trending those changing conditions. If the endpoints are static to provide a baseline for each AP type and client station device type, that baseline can be trended for deviations. This will be instrumental in keeping the thresholds for the individual conversations meaningful.

## 6. Conclusion

Aggregating quality of experience data at scale can be done multiple ways. The interesting effect WiFi has on this quality of experience is that one to many or 2 way conversation is where the first operation is computed. Measurement data from specific thresholds for endpoint pairings will be the fundamental step before rolling the data into aggregate trends. Anything that is connected to the AP that has a cataloged resource is considered an endpoint pairing, such as AP to gaming device or AP to IoT device. The basic element to measure all of the following remain consistent in WiFi quality of service: interface statistics, radio metrics and client connection statistics. Counting the amount of occurrences where an acceptable service delivery was not met for those threshold pairings is paramount.



**Figure 4 - Aggregate Trending**

Once those counts are listed, it is important to give those counts context. Using those as a percentage in the overall service delivery can setup a baseline. In the above referenced figure, thresholds are applied across conversation pairings and counted as a percentage to assess impact on the wireless customer experience. The percentage of time polling systems captured the signal and bandwidth requirement below the configured threshold is trended. The bottom chart shows the distribution for the contributing reasons why the service delivery was not satisfied. If those endpoint pairings, such as service delivery for video streaming devices are contextualized, the ability to filter on those specifics can understand the trends for capacity management can be shown at an aggregate level. If that order of operations is not met, aggregate data will remain less meaningful.

## Abbreviations

AP	access point
bps	bits per second
Hz	hertz
ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers

## Bibliography & References

Figure 1: Plashal, Andres. “Differentiating 5GHz and 2.4GHz Frequency Bands” *July 2020*

<https://andres.plashal.com/2018/blogging/differentiating-5gvs2g-frequency-bands/>

Figure 2: Metageek Training Documentation. “The Thress Main Causes of WiFi Interference” *July 2020*

<https://www.metageek.com/training/resources/why-channels-1-6-11.html>

# Developing Installation Guidelines For Wi-Fi Managed Devices

A Technical Paper prepared for SCTE•ISBE by

**Albert Garcia**  
Executive Director, TPX CPE  
Comcast  
215-286-7884  
Albert\_Garcia3@Comcast.com

# Table of Contents

<b>Title</b>	<b>Page Number</b>
1. Introduction.....	3
2. Problem Statement .....	3
3. Key Metrics.....	3
3.1. AirTime .....	3
3.1.1. Wi-Fi 6 .....	3
3.2. Received Signal Strength Indicator (RSSI).....	5
4. Prototyping .....	6
5. Characterizations .....	7
6. Baselineing.....	8
7. Testing.....	8
8. Conclusion.....	8
Abbreviations .....	9

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - AirTime Analysis Between Xfinity Camera and AP with Different Beamforming Settings .....	5
Figure 2 - Throughput Versus Frame Size Performance for Various Xfinity Managed Video Clients .....	6
Figure 3 - AP Throughput Versus Frame Size - Frame Aggregation Enabled .....	7
Figure 4 - AP Throughput Versus Frame Size - Frame Aggregation Disabled .....	7
Figure 5 - Comcast Test House. Whole Home Setup. Cross-sectional view. ....	8

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1. Xfinity Camera AirTime utilization versus RSSI analysis.....	4

# 1. Introduction

The proliferation of MSO-managed Wi-Fi devices has created a new set of challenges to define best install practices, capacity and troubleshooting. Modeling and characterizing capacity, as it pertains to multiple devices creating diverse traffic models, is challenging and evolving. Overcoming difficulty in the translation of a repeatable test lab environment, as a baseline to real world customer experience, will be discussed. This paper will focus on approaches to provide recommended device installation guidelines through a four-prong approach of prototyping, characterizing, baselining and real world testing.

This paper loosely follows the steps that Comcast took prior to the launch of the first wireless video settop. While Wi-Fi was not new, its application to a managed device was novel to the organization. Installation guidelines, much like those that existed for coaxial-connected devices, had to be developed based on new metrics which needed to be standardized across the various device combinations. Similarly, device settings needed to be adapted for this new medium through a cycle of prototyping and testing. Lastly, a new approach to whole home testing had to be taken to account for these new devices in the ecosystem.

## 2. Problem Statement

MSO-provided services over Wi-Fi create a unique set of challenges. On the radio interface, connected devices compete for resources. It is imperative to create a framework to bound the services that could be provided to a given home based on their type of service and equipment. Many Wi-Fi products are ubiquitous and have intensive bandwidth demands. Each MSO business unit looking to maximize its offerings to maximize its competitiveness – offer more devices with higher resolutions. Having a capacity model that translates to customer experience is imperative for scalable hardware and software support in a connected home. That capacity model must evolve in both terms of connected devices in a home as well as connected device traffic profiles and must be bound by installation guidelines.

## 3. Key Metrics

Through our research and testing we decided to use AirTime as the metric to quantify “how much Wi-Fi” a device or a service should be allowed in a home when installed at the recommended RSSI level.

### 3.1. AirTime

Wi-Fi began as a listen-before-talk service where only 1 device on the given channel can “talk” at a time. When a device is allowed to “talk” it will require a finite amount of time, AirTime, to complete its data transaction. The ability to transmit and/or receive data quickly and efficiently will determine how much AirTime a client will consume in a home.

#### 3.1.1. Wi-Fi 6

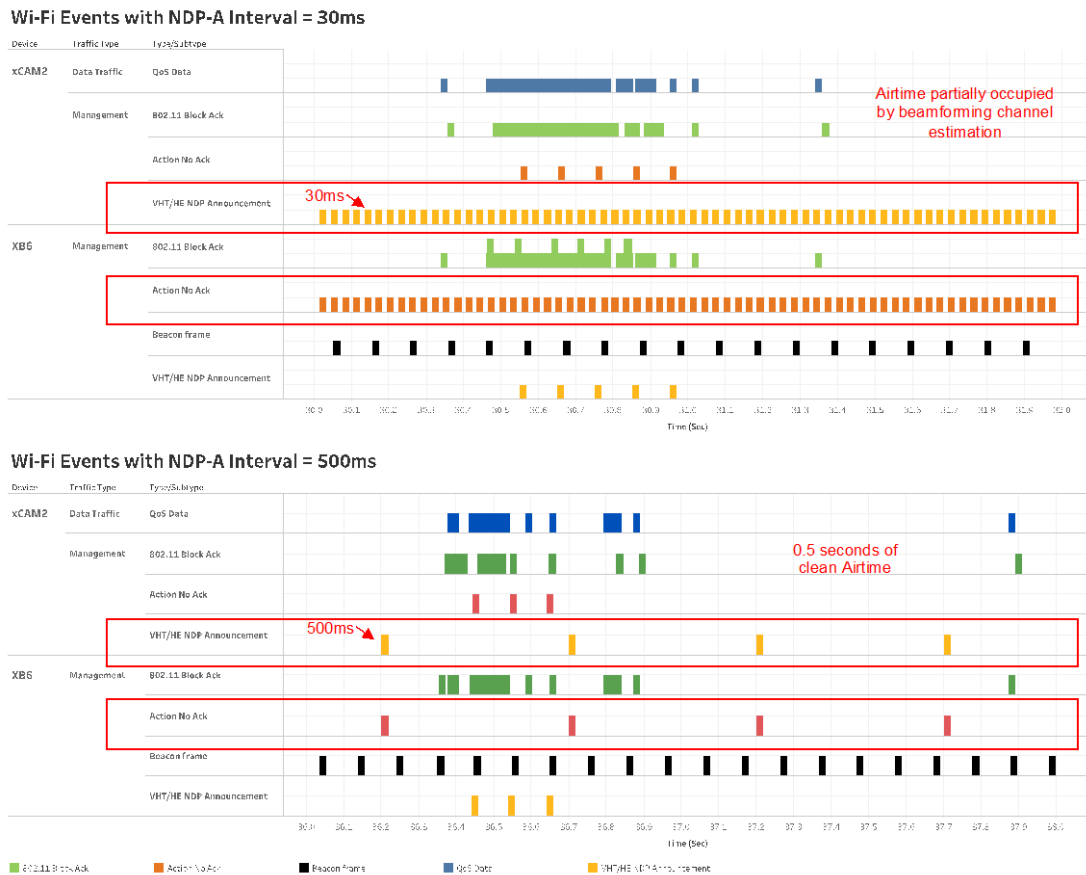
The charter for IEEE’s latest standard, Wi-Fi 6, benchmarks efficiency for high density Wi-Fi deployments. No longer was the goal to increase a given device’s bandwidth, but rather to have provisions in place to maximize the mechanisms to be as efficient as possible in high density environments. Features such as MU-MIMO and OFDMA are highly anticipated to aid in capacity demands on Wi-Fi’s radio interface. While OFDMA is not a new concept, it is new for Wi-Fi 6. This is a long awaited feature, where multiple client station devices can utilize airtime as a group, so as to be more efficient in an overcrowded environment. This will result in more throughput with less airtime.

Table 1 below shows some sample data collected from the analysis of a Wi-Fi5 Xfinity Camera in an RF chamber operating in 2.4GHz. The signal from a conducted Access Point (AP) is radiated in an RF chamber containing with various degrees of attenuation to emulate distance and/or obstructions between the AP and the device under test. This type of analysis allows us to understand the relationship between an installation closer to the AP (lower signal attenuation and higher RSSI) and the AirTime that the device will consume at that range.

**Table 1. Xfinity Camera AirTime utilization versus RSSI analysis**

Chamber Attenuation dB	RSSI (Average) dBm	SNR (Average)	Average Channel Utilization (%)	Channel Utilization 90 <sup>th</sup> Percentile (%)	Channel Utilization (%) Max	Channel Utilization (%) Min	Channel Utilization (%) STDEV	TX Rate (Max) Mb/s	TX Rate (Min) Mb/s	RX Rate (Max) Mb/s	RX Rate (Min) Mb/s
0	-56	49.28	6.26	8.1	24.7	3.1	3.09	144	122	144	78
5	-61	44.5	5.81	6.9	17.6	3.9	2.02	144	132	130	77
10	-67	38.5	5.98	7.3	11.3	4.7	1.42	144	121	129	78
15	-71	34.2	5.93	7.3	11.3	3.1	1.31	144	117	124	78
20	-68	37.3	6.72	9.0	22.3	3.5	3.42	144	133	116	77
25	-72	32.97	6.19	7.4	20.0	3.9	2.36	143	121	103	58
30	-76	28.69	6.07	7.1	16.5	3.5	2.25	125	106	89	52
35	-81	24.15	7.45	11.3	21.9	3.9	3.60	102	82	77	39
40	-86	19.86	8.60	12.3	16.0	3.5	3.08	69	56	52	29
45	-90	15.91	10.65	18.3	42.7	3.5	6.69	58	54	26	13
50	-93	15.06	13.06	23.4	41.2	5.8	8.01	52	44	26	6

Understanding not only the bandwidth requirements of the devices in a connected home, but also the time it takes on the air to get the required bandwidth, is a key metric for building a diverse capacity model. Breaking down what components are included in that measurement is pivotal. *Figure 1* below illustrates the various tasks consuming air time in the communications between an Xfinity Camera and AP and contrasts the differences between two beamforming settings. This analysis can help understand what elements, if any, beyond the data traffic are significantly contributing to the AirTime consumption of a given client. In this case, for a static device like the Xfinity Camera, the shorter beamforming interval was not only unnecessary but harmful to the overall performance of the home network by greatly increasing AirTime consumption.



**Figure 1 - AirTime Analysis Between Xfinity Camera and AP with Different Beamforming Settings**

Once the AirTime model is properly measured it will provide a foundation for install guidelines. Airtime utilization has some nuances for the measurement across SoC platforms. Often SoC vendors report the AirTime consumed by their devices differently. By including variance of metrics into their reporting such as, management overhead frames inclusive of NAV times, probe responses, beacons, SIFC intervals, the collected data may have to be normalized as a result.

### 3.2. Received Signal Strength Indicator (RSSI)


More complex than defining “how much Wi-Fi” a device should be allocated was finding an installation guideline metric that would be both easily accessible and consistent across the various System On Chip (SoC) vendors that our devices employed. This metric needed to not only be easily understood by our customers, but also accessible to our technicians, via telemetry and/or on-screen diagnostics. RSSI measured from the access point (AP) can be used to determine the given attenuation in a model and is the basis for install guidelines. RSSI at the STA could be used but it is often more difficult to use for IoT devices that are not connected to a screen. Comcast agreed to use RSSI understanding its limitations.

With an understanding of the AirTime consumption of a managed device, or devices, at a given RSSI level it is possible to begin to model the approximate theoretical Free AirTime (FAT) available to a consumer that subscribes to various services.

## 4. Prototyping

Early device prototyping became an integral phase of our characterization efforts, as it allowed us to verify our RF architectural decisions (2x2 / 3x3 / 4x4) as well as to begin the process of fine-tuning the settings of our devices firmware for an optimal user experience. Ensuring that the device hardware is capable of the theoretical maximum values provides the first step in Wi-Fi hardware verification. This is typically completed in a conducted laboratory setting. Adding attenuation to these test suites will yield valuable data points to build foundational measurements. These are usually referred to Rate vs Range (RvR) testing and Rate vs Orientation (RvO). At this prototyping phase, changing variables, such as driver settings and traffic models, is essential to qualify device behavior. Testing different protocols, such as TCP vs UDP, or modifying frame sizes, help prototype the device behavior.

Default driver settings for a Wi-Fi device can dramatically change the experience, depending on the overall desired application requirements. *Figure 2* below shows the impact of Wi-Fi frame size on the throughput capabilities of various Xfinity managed video clients. It is important to understand these impacts and bound the settings so that the devices will operate within their most efficient range.




**Figure 2 - Throughput Versus Frame Size Performance for Various Xfinity Managed Video Clients**


Another setting we found to have a significant impact on the Wi-Fi performance of a device and/or home is frame aggregation. Wi-Fi uses frame aggregation to be more efficient with regards to airtime. This can come at a price. In noisy environments, there is a higher chance for a collision or error, which results in all aggregated having to be retransmitted again. This will take longer in an already overcrowded space. Some SoCs will have algorithms in place to dynamically measure the environment and aggregate only when the conditions are satisfactory.



Figure 3 and Figure 4 below illustrate the differences in throughput for a given AP when Frame Aggregation is enabled or disabled. It is important to understand the tradeoffs of this setting for various system load levels and frame sizes.



**Figure 3 - AP Throughput Versus Frame Size - Frame Aggregation Enabled**



**Figure 4 - AP Throughput Versus Frame Size - Frame Aggregation Disabled**

## 5. Characterizations

Understanding how devices behave with the various APs and extenders is foundational for an operator to understand. Modeling can be challenging, given all the combinations and variables possible. In the prototyping phase, Wi-Fi specification compliance is mandatory. Achieving those theoretical maximum values will validate the hardware and the default driver settings for the conducted testing phase. In the IEEE Wi-Fi specifications there are also plenty of optional and other settings that can be tuned to achieve desired output, such as maximizing throughput or minimizing latency. As stated earlier, these settings can be adversely affected in noisy or crowded environments. How different settings affect the in home experience can produce different results, given the variety of different connected clients at varied RSSI levels. We instituted a rigorous and formal program to characterize each of our managed devices, (a) on its own with each of our APs and Extenders and (b) in the presence of our other managed devices. Individual characterization is performed pre-launch and frequently thereafter, as firmware, in either the managed device or APs, evolves. Whole home characterization is performed periodically in a radio silence test house such as one shown in Figure 5 where all the metrics for the various devices are monitored and compared against the individual values.



**Figure 5 - Comcast Test House. Whole Home Setup. Cross-sectional view.**

## 6. Baselining

Through the characterization of the various devices the relationships between AirTime consumption and installation range, through the measure of RSSI, become apparent. This data in combination with the business goals allows the Product and Business owners to begin to define combinations of managed clients with a given AP that strike a balance between capacity and customer experience. This somewhat iterative, and often evolving process, will create baseline device configurations that can be reliably deployed and installed.

## 7. Testing

The capacity models and installation guidelines drawn from the characterization and baselining efforts need to be corroborated in real world scenarios. This can be achieved via (a) controlled Test House tests and (b) device telemetry. A radio silent Test House can be converted into a real world scenario with the injection of noise and/or Wi-Fi traffic; device combinations can be tested against various APs with given software builds to confirm the recommended guidelines. Similarly, the real time telemetry provided by our APs can provide information to illustrate how given a device combination is performing under less controlled and more random scenarios. This test and telemetry data can be used as a feedback mechanism to further refine the characterization and baselining of the devices and software under design.

## 8. Conclusion

In spite of Wi-Fi being a live, shared medium, it is possible for an MSO to deterministically provide bounds to the number of managed devices and services based on Wi-Fi Airtime and RSSI installation guidelines.

A device design process based on prototyping, characterization, baselining and real world testing will provide product teams with the necessary information to provide capacity and installation recommendations that maximize both the number of services provided, and the customer experience.

## Abbreviations

AP	access point
bps	bits per second
IEEE	Institute of Electrical and Electronics Engineers
MIMO	Multiple In Multiple Out
MSO	Multiple Services Provider
MU-MIMO	Multi User MIMO
OFDMA	Orthogonal Frequency Domain Multiple Access
FAT	Free AirTime
RSSI	Received Signal Strength Indicator
RvO	Rate versus Orientation
RvR	Rate versus Range
SoC	System on Chip
SCTE	Society of Cable Telecommunications Engineers
Wi-Fi	Wireless Fidelity

# Tele-Everything and Its Impact to The Network

A Technical Paper prepared for SCTE•ISBE by

**Matthew Tooley**

Vice President of Broadband Technology  
NCTA – The Internet & Television Association  
Washington, DC  
(202) 222-2479  
mtooley@ncta.com

**William A. Check, Ph.D.**

Senior Vice President, Technology and Chief Technology Officer  
NCTA – The Internet & Television Association  
Washington, DC  
(202) 222-2477  
bcheck@ncta.com

**Rob Rubinovitz**

Vice President, Research and Economic Analysis  
NCTA – The Internet & Television Association  
Washington, DC  
(202) 222-2359  
rrubinovitz@ncta.com

**Jim Partridge**

Vice President, Industry and Technical Analysis  
NCTA – The Internet & Television Association  
Washington, DC  
(202) 222-2457  
jpartridge@ncta.com

# Table of Contents

Title	Page Number
1. Introduction .....	4
2. Cable Network Dashboard .....	4
3. Other Open Source Data.....	7
3.1. Mobile Wireless.....	7
3.2. Wireline .....	8
3.3. Small Cable Operators .....	9
3.4. Global Internet.....	10
3.4.1. Nokia Deepfield.....	10
3.4.2. Sandvine.....	12
3.4.3. SamKnows .....	13
3.4.4. RIPE Atlas .....	14
4. Observations.....	15
4.1. Traffic Growth.....	16
4.2. Service Delivery Infrastructure .....	16
4.3. Upstream Traffic and Telework .....	18
4.4. Upstream and Vacation Homes .....	18
4.5. Videoconferencing and Video Streaming .....	19
4.5.1. Average bandwidth per subscriber during busy-hour .....	19
4.5.2. Traffic Ratio .....	21
5. Conclusion .....	21
Abbreviations.....	22
References.....	22

## List of Figures

Title	Page Number
Figure 1 - National peak utilization growth. Source: NCTA .....	5
Figure 2 - Week-over-week change in peak utilization. Source NCTA .....	5
Figure 3 - Aggregated service group utilization. Source: NCTA.....	6
Figure 4 - Example of the state level service group peak utilization used in the NCTA Network Performance dashboard. Source: NCTA.....	7
Figure 5 - Wireless data change from baseline. Source: CTIA.....	8
Figure 6 - Wireline traffic change from baseline Adapted from USTelecom.....	8
Figure 7 - Download consumption. Source: ACA Connect.....	9
Figure 8 - Upload consumption. Source: ACA Connect .....	9
Figure 9 - Reported change in traffic patterns. Source: Nokia Deepfield .....	10
Figure 10 - Change in application usage. Source: Nokia Deepfield.....	11
Figure 11 - Videoconferencing usage growth. Source: Nokia Deepfield .....	11
Figure 12 - Change in Zoom and WebEx usage. Source Nokia Deepfield .....	12
Figure 13 - Sandvine global application category total traffic share Source: Sandvine. ....	13
Figure 14 - SamKnows measured change in download speed by U.S. state between March 12 and March 24 Source: SamKnows .....	14

Figure 15 - Round-trip time for major cable operators.....	15
Figure 16 - Round-trip times March 15-30.....	15
Figure 17 - Small cable operator transit utilization. ....	17
Figure 18 - Netflix service delivery (CDN + off-net = via peering).....	17
Figure 19 - Upstream growth and telework by state combined for the period March 1, 2020 – April 15, 2020. ....	18
Figure 20 - Seasonal homes and utilization combined for the period March 1, 2020 – April 15, 2020.....	19

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Aggregated service group utilization.....	6
Table 2 - Traffic engineering calculations for cable networks.....	16
Table 3 - Average Subscriber Bandwidth During Peak Busy Hours. ....	20
Table 4 - Videoconferencing bandwidth requirements. ....	20
Table 5 - Video bandwidth requirements. ....	21

# 1. Introduction

The COVID-19 (2019 novel coronavirus) pandemic caused governments around the world to issue shutdown orders resulting in a sudden shift in network traffic patterns as subscribers started using their broadband connections for tele-everything – working from home, remote learning, entertainment, social interactions, and commerce. As part of the crisis management effort, and to keep the public and various levels of government informed, operators posted reports of the status of their broadband networks in handling the increased usage.

To develop these reports, cable operators shared aggregated anonymized data with NCTA – The Internet & Television Association. In this paper, we first look at the data that the cable operators shared with NCTA as well as similar data from wireline and wireless operators. We then cross-reference the cable data with other third-party data to make some additional observations on the performance of the broadband access networks.

## 2. Cable Network Dashboard

NCTA worked with nine of the U.S.’s leading cable companies<sup>1</sup> to collect and aggregate anonymized network performance data for a publicly accessible network performance dashboard (NCTA, 2020). The dashboard reported on two metrics:

1. change in peak utilization
2. service group utilization grouped by state


The dashboard went live on April 4, 2020 and was still active at the time of the writing of this paper in August 2020. The data set for the dashboard includes data going back to March 1, 2020 – prior to when states started issuing shutdown orders. This was in order to measure and report on the impact of the orders on the network. The dashboard also provided insights into what happened when states started re-opening in late April 2020.

Change in peak utilization was chosen as the primary metric for the dashboard because networks are engineered for peak capacity, and this single metric provided the best snapshot of how well the networks were performing. Figure 1 shows the overall change in the peak utilization for both the upstream (US) and downstream (DS) traffic going back to March 1, 2020 to show the change from pre-shutdown levels. It is important to note that the change in utilization is not directly analogous to the change in consumption because the overall capacity of the networks are fluid; operators could be adding capacity each week. Even though a change in utilization is not the same as consumption, it still provides a good proxy for how data consumption changed as a result of the shutdown.

Peak upstream utilization occurred the week ending April 18, 2020, with a 35% increase over the pre-shutdown levels. For downstream traffic, the largest change in peak utilization occurred for the week ending March 28, 2020, at 20%.


---

<sup>1</sup> Companies who shared data included: Altice, CableOne, Charter, Comcast, Cox, GCI, Mediacom, Midco, and Sjoberg’s.



**Figure 1 - National peak utilization growth. Source: NCTA**

Figure 2 shows the week-over-week change in the peak utilization. The largest week-over-week change for the upstream occurred the week ending March 28, 2020, with a 7.3% increase and for the downstream the week of June 13, 2020, with a 7.8% increase. Again, it is important to note that this is the relative utilization of available capacity for that week.



**Figure 2 - Week-over-week change in peak utilization. Source: NCTA**


In addition to the aggregated change in peak utilization at the national level, MAC (medium access control) domain service groups utilization were collected. A DOCSIS (Data Over Cable System Interface Specification) service group is either a MAC domain upstream service group or a MAC domain downstream service group. A MAC domain downstream service group refers to the set of downstream channels from a port on a CMTS (cable modem termination system) line card that reach a fiber node and a MAC domain upstream service group refers to the set of upstream channels from the same MAC domain that are reached by a single cable modem. Operators shared the peak utilization levels for each service group and further grouped these by utilization level (<85%, 85%-95%, 95%-99%, and >99%) as shown in Table 1 and in Figure 3



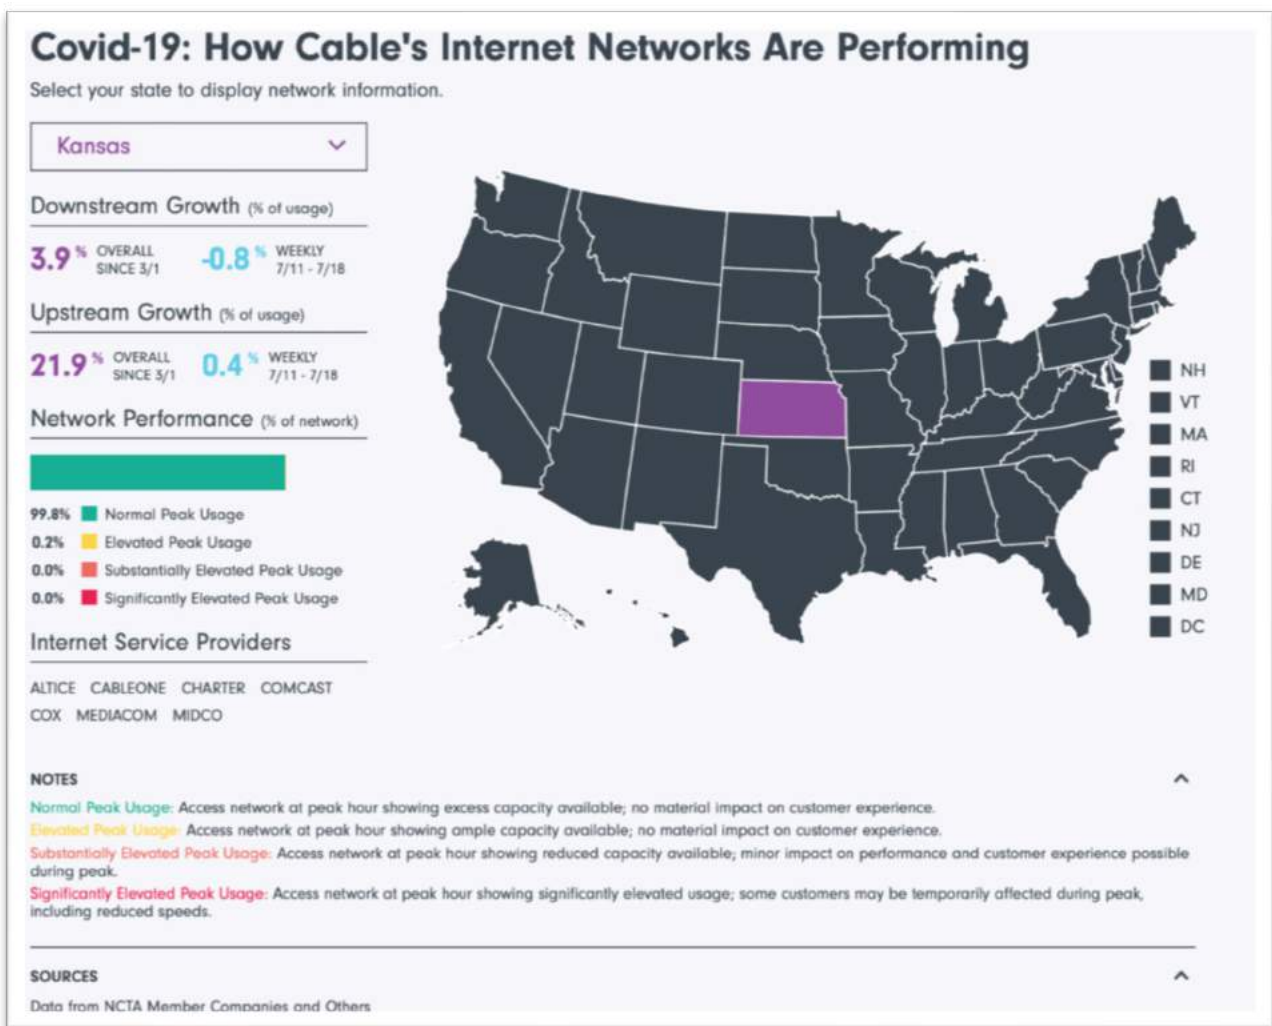
Less than 0.35% of the service groups were ever above 95% utilization. For the dashboard, the service groups were aggregated by state to show how well the DOCSIS network was operating on a regional basis as shown in Figure 4.

**Table 1 - Aggregated service group utilization.**

<i>Week</i>	Service Group Utilization Level			
	< 85%	85% - 95%	95% - 99%	>99%
14	95.8%	3.9%	0.29%	0.04%
15	95.8%	3.8%	0.29%	0.06%
16	96.3%	3.5%	0.22%	0.02%
17	96.4%	3.4%	0.16%	0.01%
18	96.0%	3.7%	0.23%	0.03%
19	97.1%	2.7%	0.12%	0.05%
20	97.1%	2.8%	0.13%	0.01%
21	97.0%	2.8%	0.16%	0.02%
22	97.9%	2.0%	0.09%	0.00%
23	98.1%	1.8%	0.07%	0.02%
24	97.7%	2.1%	0.13%	0.03%
25	98.3%	1.6%	0.07%	0.01%
26	98.5%	1.4%	0.05%	0.02%
27	98.3%	1.6%	0.08%	0.01%



**Figure 3 - Aggregated service group utilization. Source: NCTA**




**Figure 4 - Example of the state level service group peak utilization used in the NCTA Network Performance dashboard. Source: NCTA**

### 3. Other Open Source Data

The mobile wireless and wireline carriers, along with small cable operators, also shared network performance data via their respective trade associations.

#### 3.1. Mobile Wireless


CTIA reported on the cellular network performance (CTIA, 2020). Mobile data use peaked at 22.6% above the pre-shutdown average on June 22, 2020.



**Figure 5 - Wireless data change from baseline. Source: CTIA**

### 3.2. Wireline

USTelecom reported on the wireline (DSL, fiber) network performance by reporting on the change in traffic (aggregated upstream & downstream) compared to the pre-shutdown level (USTelecom, 2020). The peak traffic change occurred the week of April 16, 2020, at 27% higher than the baseline, and then trended down to about 12%.



**Figure 6 - Wireline traffic change from baseline Adapted from USTelecom**

### 3.3. Small Cable Operators

ACA Connects, the trade association representing many of the smaller U.S. cable companies, reported that peak download usage was up 24% and peak upload usage was up 34% for its members (ACA Connects, 2020). Download consumption peaked at 44.7% (relative to February 1, 2020) on March 27, 2020, and the upload consumption peaked at 68.8% on March 24, 2020, as shown in Figure 7 and Figure 8.




Figure 7 - Download consumption. Source: ACA Connect





Figure 8 - Upload consumption. Source: ACA Connect

### 3.4. Global Internet

#### 3.4.1. Nokia Deepfield

At NANOG 79 (North American Network Operators' Group) in June 2020, Craig Labovitz of Nokia Deepfield gave a presentation on the impact of the pandemic on the internet on a global scale, in which he discussed impacts on networks both domestically and internationally (Labovitz, 2020). As shown in Figure 9, Labovitz reported that most networks typically have 40-50% traffic growth in one year, and many saw a 45% change during the first four weeks of the shutdown.



**Figure 9 - Reported change in traffic patterns. Source: Nokia Deepfield**

Labovitz also reported on the change in application usage as shown in Figure 10 after the first week of the shutdown, with the secure messaging app “WhatsApp” and videoconferencing apps in general showing the largest growth in usage. Figure 11 shows the growth in videoconferencing with Zoom usage growing as much as 700% on select networks. Figure 11 also shows how Zoom moved to a multi-CDN (content delivery network) strategy from a single CDN delivery strategy. Figure 12 illustrates how WebEx had somewhat consistent growth for each day of the week, while Zoom experienced its highest growth on Saturday and Sunday.




Figure 10 - Change in application usage. Source: Nokia Deepfield

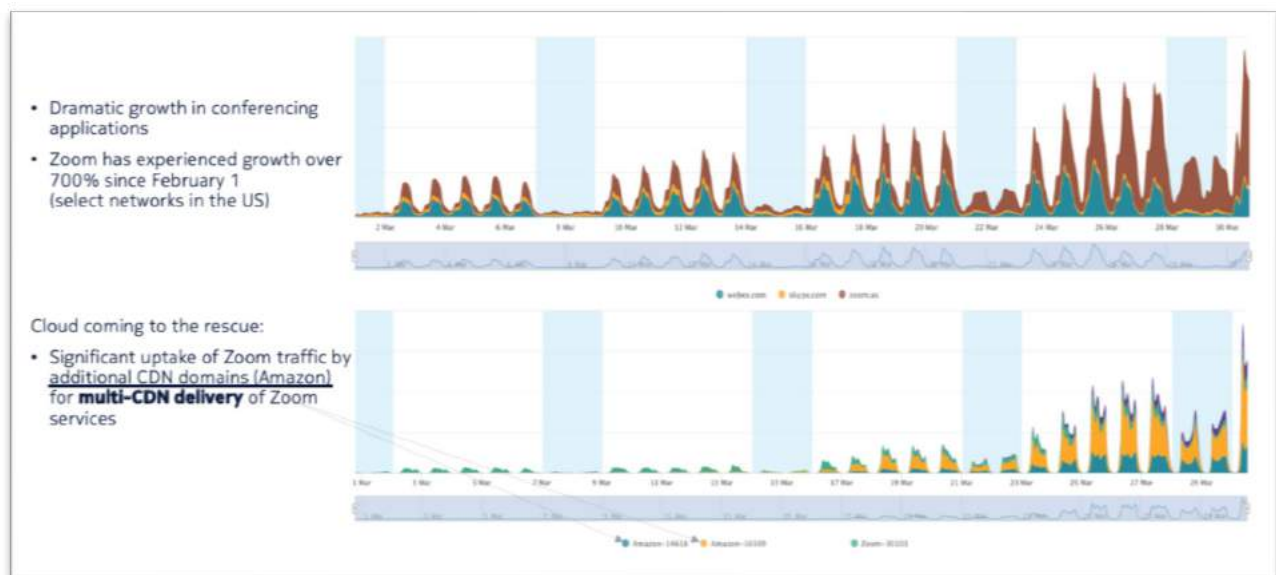



Figure 11 - Videoconferencing usage growth. Source: Nokia Deepfield



**Figure 12 - Change in Zoom and WebEx usage. Source Nokia Deepfield**

### 3.4.2. Sandvine

Sandvine reported that Internet traffic grew almost 40% due to the shutdowns (Sandvine, 2020). In addition, Sandvine also reported on the total traffic by application category during the shutdown as shown in Figure 13, with video streaming having the largest share of the traffic at 57.6% and growing 2.2%.



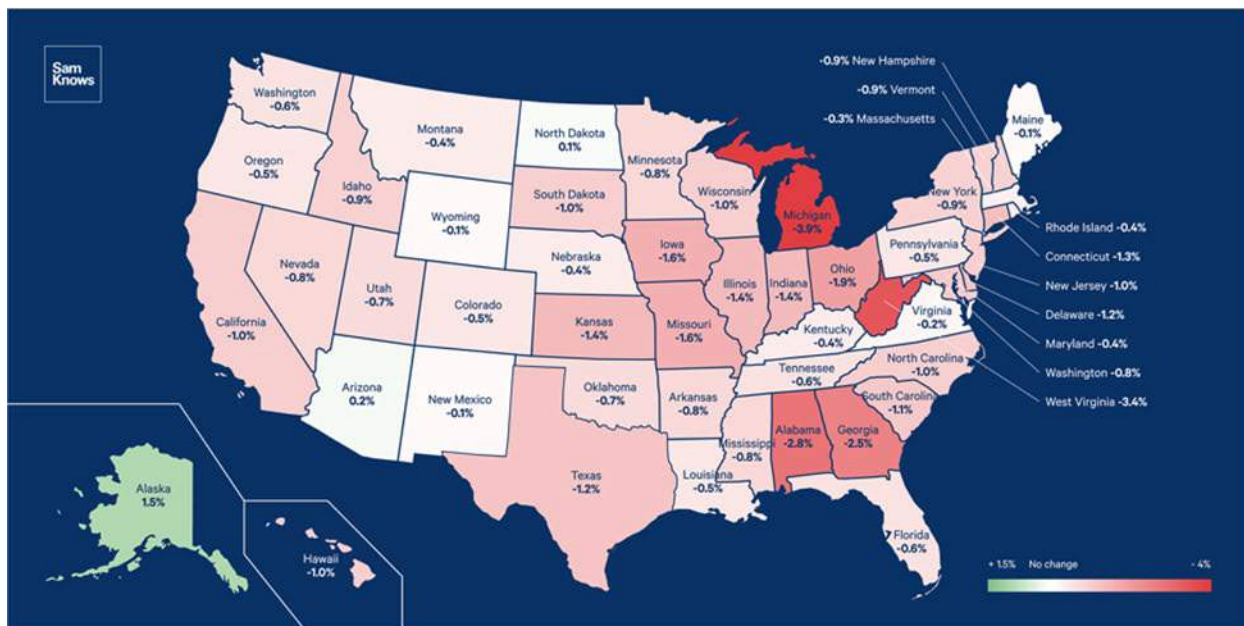


**Figure 13 - Sandvine global application category total traffic share Source: Sandvine.**

### **3.4.3. SamKnows**

SamKnows operates a large scale measurement platform that is used by ISPs, governments, academics and consumers. SamKnows gathered the test results from 500,000 homes that have a SamKnows-enabled router installed. These home routers run regularly scheduled speed tests to a major U.S.-based CDN every 10 seconds using multiple TCP (transmission control protocol) connections in parallel. SamKnows posted that most of the states only saw about a 1% decline in the download speeds as shown in Figure 14 (SamKnows, 2020). SamKnows did not provide any further explanation for what may have caused the 1% slowdown. Further it is unclear whether this reduction is within the margin of error or whether in fact the networks actually experienced a 1% decrease in downstream speed.






Based upon a sample of more than 500k homes running automated download speed tests on a regular basis. Each speed test uses 16 concurrent TCP sessions and measures to a major US CDN. Measurements from 2020-03-12 were compared against measurements from 2020-03-24 to create this comparison.

**Figure 14 - SamKnows measured change in download speed by U.S. state between March 12 and March 24 Source: SamKnows**

#### 3.4.4. RIPE Atlas


The RIPE Network Coordination Centre operates a large Internet measurement platform, RIPE Atlas, that employs a global network of probes that measure Internet connectivity and reachability. Each RIPE Atlas probe runs a set of built-in tests 24x7. One of the tests is for each RIPE Atlas probe to issue a set of pings every 15 minutes to each of the 13 root servers. The ping test results, in aggregate, can be used to monitor the latency on a network and as a proxy for detecting quality of service issues due to congestion or a change in routes.

To measure the impact of tele-everything, results were pulled from RIPE for the Atlas IPv4 ping test to the E-Root Server every 4 minutes for ten random Atlas probes on four of the largest U.S. cable networks for the period of February 1, 2020 to July 1, 2020. We grouped them by network operator to look at how the round-trip time changed from pre-shutdown to when the shutdown occurred. Figure 15 shows the round-trip times on an hourly basis for four of the major cable operators. For the most part, the round-trip time did not change very much between the pre-shutdown (February 2020) and when the shutdowns started (March 2020). During the first two weeks of the shutdown (March 15-30) we can see that there are periods when the round-trip times increase. Several of the spikes in the round-trip time occur during the early hours of the morning. Figure 16, shows the round-trip times for the last two weeks in March and we can see some spikes that could be correlated with a shift to tele-everything.



RIPE Atlas Probes Round Trip Time from Cable Networks to E.root-servers.net. E-root-servers.net uses IP Anycast to provider service from the closest location. All times shown are in GMT.

**Figure 15 - Round-trip time for major cable operators.**



**Figure 16 - Round-trip times March 15-30.**

## 4. Observations

In addition to the published data above, we can glean some additional insights from other data that cable operators and others reported about network performance.

## 4.1. Traffic Growth

In general, overall broadband usage increased due to the shutdown. The NCTA Network Performance dashboard reported an increase in the peak utilization and others reported overall traffic growth between 20-50% (CTIA, 2020); (USTelecom, 2020); (ACA Connects, 2020); (Labovitz, 2020); (Sandvine, 2020)).

We can get an approximation of the magnitude of the traffic growth on the cable networks by comparing the forecasted growth of peak downstream and upstream bandwidth per subscriber with the observed growth in the peak utilization at the start of the shutdown. Ulm and Cloonan reported that in 2019 the downstream average bandwidth per subscriber during busy-hour (DS Tavg) was 1.97 megabit per second (Mbps) with a 5 year average CAGR (compound annual growth rate) of 37.8% and the upstream average bandwidth per subscriber (US Tavg) was 140 kbp/s with a 5 year average CAGR of 18.8% (Ulm & Cloonan, 2019). Using these numbers the DS Tavg is forecasted to grow about 3% each month and the US Tavg forecasted to grow about 1.5% per month on average. If we compare these with the observed change in the peak utilization on the cable networks of 20% in the downstream and 35% in the upstream, it becomes clear that traffic grew substantially as a result of the shutdown.

**Table 2 - Traffic engineering calculations for cable networks.**

	Downstream	Upstream
<b>Tavg</b>	1.97 Mbps	140 kbp/s
<b>5-year CAGR</b>	37.8%	18.8%
<b>2020 forecasted monthly growth</b>	3%	1.5%
<b>2020 change in peak utilization</b>	20%	35%

In addition to traffic growth, we can also infer that the shift to tele-everything resulted in cable networks experiencing about six months to a years' worth of traffic growth in a month and that this was consistent with what other ISPs around the globe experienced. We can also infer from the SamKnows download speed measurements that this growth did not have a measurable impact on the subscribers download or upload speeds. And finally, the shift to tele-everything caused a step-function in demand for bandwidth. It remains to be seen whether this is a temporary step function that will return to the pre-shutdown levels at some point, or if this is a permanent shift.

## 4.2. Service Delivery Infrastructure

ACA Connects posted peak transit network capacity utilized by day for its larger members as shown in Figure 17. We can see that the transit utilization did not change with the shift to tele-everything and we can infer from this that the service delivery infrastructure of content delivery networks, caches, and peering agreements with content partners worked well. This is consistent with a similar report by Labovitz, as shown in Figure 18 on the performance of the Netflix service delivery system where before the shutdowns Netflix had a 63% cache hit rate, and 46% cache hit rate once the shutdowns began (Labovitz, 2020).




Figure 17 - Small cable operator transit utilization.

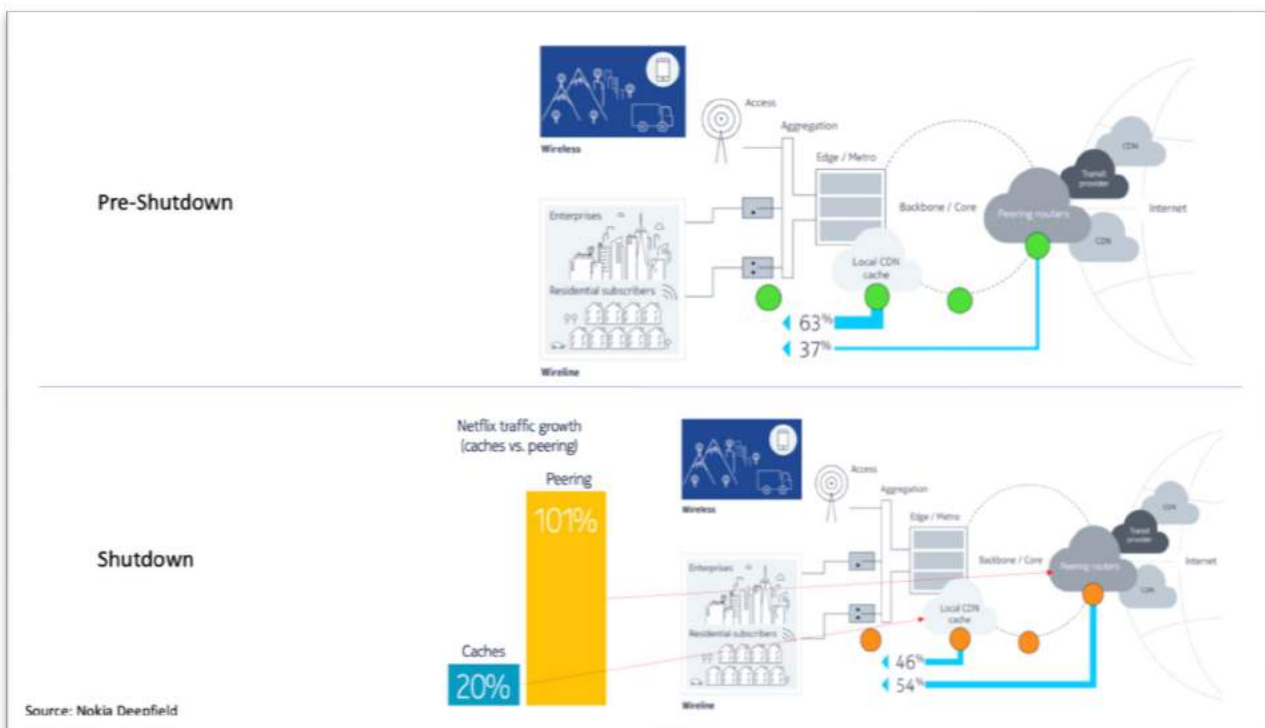
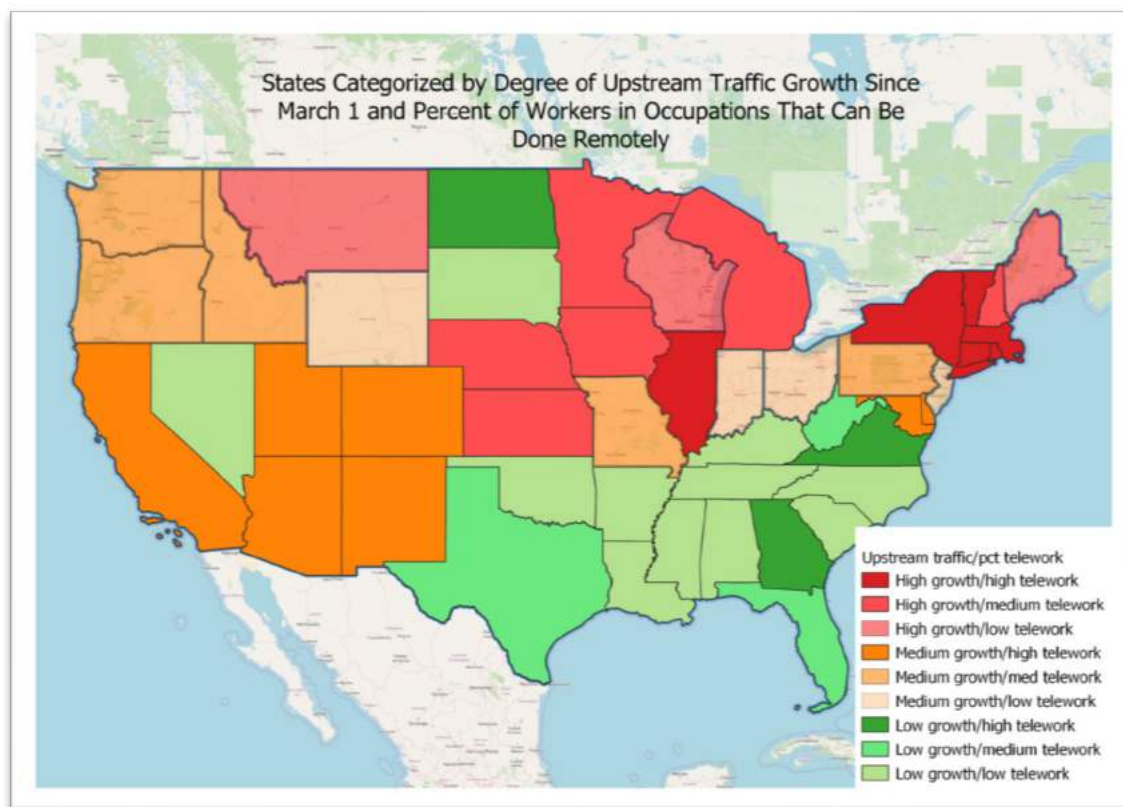


Figure 18 - Netflix service delivery (CDN + off-net = via peering).

### 4.3. Upstream Traffic and Telework

One of the big questions at the beginning of the shutdown was how would tele-everything impact the upstream traffic in the networks. We analyzed the data and attempted to correlate it with other data such as census data to see if geographic areas with high concentrations of occupations that can be done remotely might explain changes in upstream traffic. We cross-referenced service group utilization data with census data to generate the heat map shown in Figure 19. The heat map shows the states with the largest changes in upstream utilization combined with the largest percentage of the occupations that can be done remotely. The map was generated using data from March 1 and April 15, when most of the cases were still in the state of New York, and illustrates that there is some correlation between upstream utilization growth and tele-work. As much as there is high interest in such correlations, such conclusions are tentative at best. Time of year, weather, school and other factors can influence this. However, the heat map does seem to show some correlation.



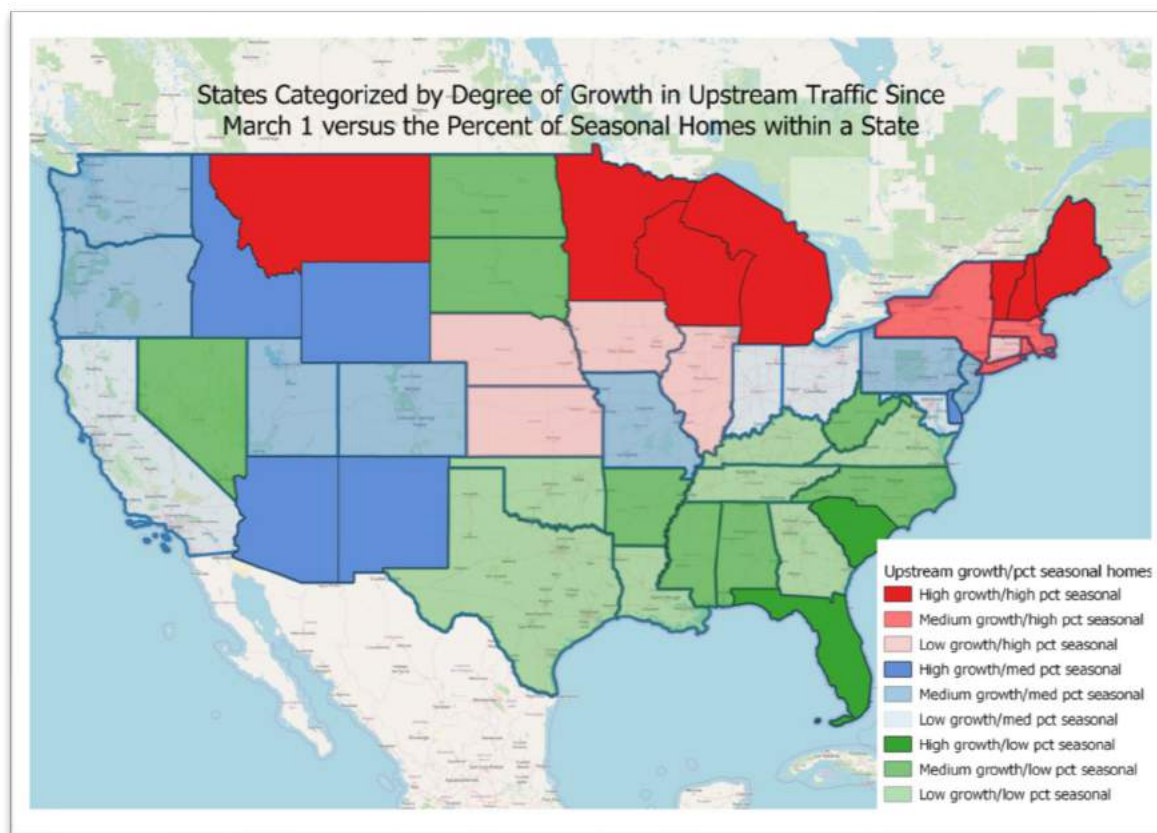
**Figure 19 - Upstream growth and telework by state combined for the period March 1, 2020 – April 15, 2020.**

### 4.4. Upstream and Vacation Homes

Similarly, we also looked to see if there was any potential correlation with areas experiencing large changes in upstream utilization due to subscribers migrating to their vacation homes to wait out the shutdown. Figure 20 shows utilization combined with vacation home data. The northern states with a large percentage of vacation homes experienced the largest increases in upstream peak utilization. This may be because many of the vacation homes in the north are three-season homes (spring, summer, and



fall) and therefore, what we may be observing is vacation homeowners moving to their homes earlier than normal as part of the shelter in place orders. Again, these types of correlation are suppositional.



**Figure 20 - Seasonal homes and utilization combined for the period March 1, 2020 – April 15, 2020.**

## 4.5. Videoconferencing and Video Streaming

Another concern raised was the potential impact on upstream traffic with a high usage of videoconferencing for work, school, and entertainment. It was reported that videoconferencing applications such as Zoom, with symmetrical bandwidth requirements, experienced as much as a 700% growth with the largest growth occurring on weekends.

### 4.5.1. Average bandwidth per subscriber during busy-hour

Using the average bandwidth per subscriber during busy-hour, Tavg, for the downstream and upstream from 2019 and the reported 5-year average CAGRs (Ulm & Cloonan, 2019), we can estimate the forecasted 2020 Tavg for the downstream and upstream around the time of the shutdown when consumers switched to tele-everything. In Table 3 we show the calculations for the forecasted Tavg and then extrapolate the change to Tavg as a result of tele-everything. We estimate that the downstream Tavg grew by 430 kbp/s, growing to 2,582 kbp/s, and the upstream Tavg grew by 51 kbp/s to 198 kbp/s.

**Table 3 - Average Subscriber Bandwidth During Peak Busy Hours.**

	Description	Formula	Down	Up
(A)	2019 average bandwidth/sub during busy hour (Tavg) <sup>2</sup> , kbp/s	A[DOWN] / A[UP]	1966	140
	2019 downstream-to-upstream traffic ratio		14	
(B)	2019 CAGR <sup>3</sup>		37.80 %	18.80 %
(C)	Estimated monthly growth in Tavg	B/12	3.15%	1.57%
(D)	Estimated growth in Tavg for Jan-Mar 2020, kbp/s	A*C*3	186	7
(E)	Estimated Tavg at beginning of April 2020	A+D	2152	147
(F)	Estimated DS:US traffic ratio for April 2020	E[DOWN] / E[UP]	15	
(G)	Change in peak due to shutdown <sup>4</sup>		20%	35 %
(H)	Estimated growth in Tavg with shift to tele-everything, kbp/s	G*E	430	51
(I)	Estimated tele-everything Tavg, kbp/s	E+H	2582	198
(J)	Estimated tele-everything downstream-to-upstream traffic ratio	I[DOWN] / I[UP]	13	

In relative terms the upstream's peak grew more than downstream. Converting the peak utilization to an absolute number allows us to take into account the asymmetrical nature of DOCSIS broadband connections and the fact that the upstream grew from a lower base. When we do this, we can see that in absolute terms that the downstream bandwidth usage grew ~9x the upstream.

Even though the usage of videoconferencing grew much more than the usage of video streaming, the fact that video streaming uses on average about five-times more bandwidth than videoconferencing, as shown in Table 4 and Table 5, contributed to the 9x growth in absolute bandwidth in the downstream.

**Table 4 - Videoconferencing bandwidth requirements.**

	Upstream	Downstream	Average
<b>Zoom (Zoom, 2020)</b>	600 kbp/s – 1.8 Mbp/s	600 kbp/s – 1.8 Mbp/s	1/1 Mbp/s
<b>WebEx (Cisco Webex, 2020)</b>	500 kbp/s – 3.0 Mbp/s	500 kbp/s – 3.0 Mbp/s	1/1 Mbp/s
<b>MS Teams (Microsoft, 2020)</b>	500 kbp/s – 1.5 Mbp/s	500 kbp/s – 1.5 Mbp/s	1/1 Mbp/s

<sup>2</sup> J. Ulm and T. Cloonan, "The Broadband Network Evolution Continues - How Do We Get to Cable 10G?," p. 8-9

<sup>3</sup> Ibid, p. 8-9

<sup>4</sup> NCTA, "COVID-19: How Cable's Internet Networks are Performing,"

**Table 5 - Video bandwidth requirements.**

<b>Streaming Service</b>	<b>Downstream</b>
Netflix	5 Mbp/s
Hulu	3 Mbp/s
Amazon Prime	5 Mbp/s
YouTube	7 Mbp/s

#### **4.5.2. Traffic Ratio**

The shift to tele-everything had minimal impact on the downstream-to-upstream traffic ratio. In 2019, the busy-hour downstream-to-upstream traffic ratio was 14:1. We estimate that it grew to 15:1 in 2020 prior to the shutdown and with the shift to tele-everything that it declined to 13:1. This is consistent with overall traffic ratio, as reported by OpenVault, which reported that it went from 20:1 down to 16:1 (OpenVault, 2020). The reduction in the overall traffic ratio can most likely be attributed to the increased use of videoconferencing during the daytime hours for tele-work and tele-school, while the smaller change during the peak busy hours is likely due to any increased upstream usage being offset by increased video streaming in the downstream.

## **5. Conclusion**

The COVID-19 pandemic of 2020 caused a sudden shift to tele-everything that resulted in large changes in traffic patterns on the Internet and cable networks. In this paper, we looked at data collected and posted by the cable industry as well as data posted by other third parties on the performance of cable networks and the Internet in general with the shift to tele-everything. We observed that the shift to tele-everything caused internet traffic to grow 30-50% and that traffic on the cable networks grew 20-35%, with the traffic growth on cable networks somewhat independent of the size of the operator. We observed that the service delivery infrastructure of content delivery networks and caches worked well, with most of the traffic growth being observed in the access networks and not on the transit links. We also observed that even with the sudden growth in traffic, there was no measurable impact on the end user's quality of experience, as measured download speeds only declined about 1% and there was no measurable change in the round-trip time for the RIPE Atlas test probes.

We also looked at the impact of the shift to telework and the early opening of northern vacation homes on the upstream usage. We observed that there was a loose correlation between the two.

And finally, we looked at the impact of the growth in videoconferencing and video streaming on the network. We observed that even though the upstream in cable networks in relative terms grew more than the downstream, that the absolute bandwidth growth on the downstream was 9x the bandwidth growth in the upstream and that this is because video streaming uses about 5 Mbp/s per high-definition video stream compared the 1 Mbp/s in each direction used by the videoconferencing applications. The impact to the downstream-to-upstream traffic ratio was small as the growth of video streaming in the downstream was somewhat offset by the growth in video conferencing in the upstream.

Overall, U.S. cable networks have performed well during the pandemic.



## Abbreviations

CAGR	compound annual growth rate
CDN	content delivery network
CMTS	cable modem termination system
COVID-19	corona virus disease 2019
DOCSIS	Data Over Cable System Interface Specification
DS	downstream
DSL	digital subscriber line
FTTH	fiber to the home
HFC	hybrid fiber coax
ISP	Internet service Provider
kbp/s	kilobit per second
MAC	medium access control
Mbp/s	megabit per second
NANOG	North American Network Operators' Group
SCTE	Society of Cable Telecommunications Engineers
Tavg	Average bandwidth consumed by a subscriber during the busy-hour
TCP	transmission control protocol
US	upstream

## References

- ACA Connects. (2020, July 17). *Network Performance During the COVID-19 Crisis*. Retrieved July 17, 2020, from <https://acaconnects.org/covid-19/broadband-dashboard/>
- Cisco Webex. (2020, July 22). *Bandwidth Requirements*. Retrieved July 22, 2020, from <https://help.webex.com/en-us/WBX22158/What-are-the-Minimum-Bandwidth-Requirements-for-Sending-and-Receiving-Video-in-Cisco-Webex-Meetings>
- Cloonan, T., & Ulm, J. (2017). *Traffic Engineering in a Fiber Deep Gigabit World*. 2017 SCTE Fall Technical Forum.
- CTIA. (2020, July 17). *The Wireless Industry Responds to COVID-19*. Retrieved July 17, 2020, from <https://www.ctia.org/homepage/covid-19>

- Labovitz, C. (2020, June 1). *NANOG 79*. Retrieved July 17, 2020, from [https://storage.googleapis.com/site-media-prod/meetings/NANOG79/2208/20200601\\_Labovitz\\_Effects\\_Of\\_Covid-19\\_v1.pdf](https://storage.googleapis.com/site-media-prod/meetings/NANOG79/2208/20200601_Labovitz_Effects_Of_Covid-19_v1.pdf)
- Microsoft. (2020, July 8). *Prepare Your Organization for Microsoft Teams*. Retrieved 2020 July, from <https://docs.microsoft.com/en-us/microsoftteams/prepare-network>
- NCTA. (2020, July 17). *COVID-19: How Cable's Internet Networks are Performing*. Retrieved July 17, 2020, from <https://www.ncta.com/COVIDdashboard>
- OpenVault. (2020, April 7). *COVID-19 Broadband Usage "Reaching a plateau," says OpenVault*. Retrieved July 17, 2020, from <https://openvault.com/covdi-19-broadband-usage-reaching-a-plateau-says-openvault>
- SamKnows. (2020, April 14). *SamKnows Critical Services Report: Fixed Speed (USA)*. Retrieved July 17, 2020, from <https://samknows.com/blog/samknows-critical-services-report-fixed-speed-usa>
- Sandvine. (2020). *The Global Internet Phenomena Report - COVID-19 Spotlight May 2020*. Sandvine.
- Ulm, J., & Cloonan, T. (2019, September 30). *The Broadband Network Evolution Continues - How Do We Get to Cable 10G?* Retrieved July 17, 2020, from <https://www.nctatechnicalpapers.com/Paper/2019/2019-the-broadband-network-evolution-continues>
- USTelecom. (2020, July 17). *Network Performance*. Retrieved July 17, 2020, from <https://www.ustelecom.org/research/network-performance-data/#>
- Verizon. (2020, June 11). *Verizon delivers network reliability during COVID-19 while accelerating 5G deployments*. Retrieved July 17, 2020, from <https://www.verizon.com/about/news/how-americans-are-spending-their-time-temporary-new-normal>
- Zoom. (2020, July 22). *Zoom System Requirements for Windows, macOS, and Linux*. Retrieved July 22, 2020, from <https://support.zoom.us/hc/en-su/articales/201362023-System-requirements-for-Windows-macOS-and-Linux>

# **Wireless Access Network Strategies**

## **Lessons Learned On 3.5 GHz CBRS Network Trials**

A Technical Paper prepared for SCTE•ISBE by

**Haider Syed**

Sr. Director, Wireless Engineering

Charter Communications, Inc.

6360 S Fiddlers Green Circle, Greenwood Village, CO 80111

+1 973-842-1012

Haider.Syed@charter.com

# Table of Contents

Title	Page Number
1. Introduction.....	5
2. Utilization of CBRS Shared Spectrum .....	5
2.1. Citizens Broadband Radio Service .....	5
2.2. A novel three-tier paradigm, an industry first with CBRS.....	5
2.3. Spectrum Access System .....	7
2.4. Environmental Sensing Capability .....	8
2.5. Incumbents.....	8
2.6. Protection of Federal Incumbents .....	8
2.7. Protection of Non-Federal Incumbents .....	9
2.8. Commercial SAS Trial Observations in New York City, NY .....	9
2.9. Conclusion.....	9
3. Implications to the deployment strategies .....	10
3.1. Attached Mount Deployment (Rooftop).....	10
3.2. Strand Mount Deployment .....	10
3.3. Small Medium Business (SMB) Deployment .....	11
4. Performance Characteristics of CBRS LTE Network.....	14
4.1. Characteristics of CBRS Wave .....	14
4.2. mmWave Comparison.....	15
4.3. 5G CBRS vs 4G CBRS .....	16
4.4. CBRS Coverage Reliability from Field Tests .....	17
4.5. Factor affecting throughput in CBRS network – Field Assessment .....	18
4.6. LTE Performance KPIs Observation from Field Trials.....	20
4.7. Capacity Testing .....	21
4.8. Busy Hour Traffic Analysis .....	23
4.9. Conclusion.....	25
5. CBRS Radio Devices and RF Design of CBRS Network.....	25
5.1. CBRS Radio Devices .....	25
5.2. CPE for Fixed Wireless Access .....	25
5.3. RF Design .....	26
5.4. Conclusion.....	26
6. Introduction of 3GPP Virtual RAN Split Options .....	26
6.1. Split Options .....	27
6.2. Split Options and Latency Requirements.....	28
6.3. Split Option 2 vs Option 7-2 .....	29
6.4. Conclusion.....	30
7. DOCSIS Backhaul.....	30
7.1. Low Latency Xhaul (LLX) .....	30
7.2. Low Latency DOCSIS (LLD) .....	31
7.3. Conclusion.....	31
8. Timing and Synchronization.....	32
8.1. DOCSIS Timing Protocol (DTP).....	33
8.2. Assisted GPS .....	34
8.3. Conclusion.....	34
9. Wrap-Up .....	35
Abbreviations .....	36
Bibliography & References.....	37

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 CBRS 3.5 GHz Frequency Band .....	5
Figure 2 3-Tier CBRS Structure .....	6
Figure 3 Spectrum Sharing Ecosystem .....	7
Figure 4 Attached Mount (Rooftop) Outdoor Deployment .....	10
Figure 5 Strand Mount Outdoor Deployment .....	11
Figure 6 Indoor CBSD Deployment in SMB .....	12
Figure 7 Inside-Out Deployment in Spectrum Stores .....	12
Figure 8 Outside-In Test Locations .....	13
Figure 9 CBRS 3.5GHz Antenna Properties .....	14
Figure 10 Dense Urban Morphology .....	15
Figure 11 5G CBRS vs 5G 28 GHz mmWave Results .....	16
Figure 12 Field Drive Test Plot .....	17
Figure 13 Field Single Site Performance Stats .....	18
Figure 14 Field Cluster DL Performance Stats .....	19
Figure 15 Field Cluster UL Performance Stats .....	19
Figure 16 Attached Mount Cluster KPIs .....	20
Figure 17 Cluster Drive RSRP and Throughput .....	21
Figure 18 RSRP to Throughput Function .....	21
Figure 19 Load Test Setup on Field .....	22
Figure 20 Load Test DL Performance Stats .....	22
Figure 21 Protocol Split Options .....	27
Figure 22 Sub-splits for Split Option 7 .....	28
Figure 23 Sub-splits Option 7-2 vs Split Option 2 .....	28
Figure 24 Split Option 2 .....	29
Figure 25 Split Option 2 over DOCSIS .....	29
Figure 26 Split Option 7-2 over Fiber .....	30
Figure 27 Low Latency Xhaul .....	30
Figure 28 Bandwidth Report OFF/ON .....	31
Figure 29 DOCSIS Timing Protocol .....	34

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 Inside-out Testing .....	12
Table 2 Outside-In Testing .....	13
Table 3 5G CBRS vs 5G 28 GHz mmWave Assumptions .....	15
Table 4 LTE vs NR Comparison .....	16
Table 5 Cluster Performance KPIs .....	20
Table 6 DL Load Testing .....	23
Table 7 Load Effect on Cell Radius .....	23
Table 8 Assumptions for Busyhour (BH) effect on Throughput .....	23

Table 9 Busyhour (BH) Traffic Analysis - Strand Mount .....	24
Table 10 Busyhour (BH) Traffic Analysis - Attached Mount .....	24
Table 11 Busyhour (BH) Traffic Analysis - SMB .....	24
Table 12 CPE Category .....	25
Table 13 Frequency and Phase Synchronization Requirements.....	32
Table 14 Synchronization Types - Advantages and Disadvantages .....	33

## 1. Introduction


Considering adding a wireless footprint to your network? Hear about the lessons which we have learned over the past three years of CBRS network researches, trials and how we are building a wireless network. What are the key challenges that have been overcome and the areas that need to be considered? As networks converge, wireless technology is becoming another tool to provide additional services and expand the network.

## 2. Utilization of CBRS Shared Spectrum

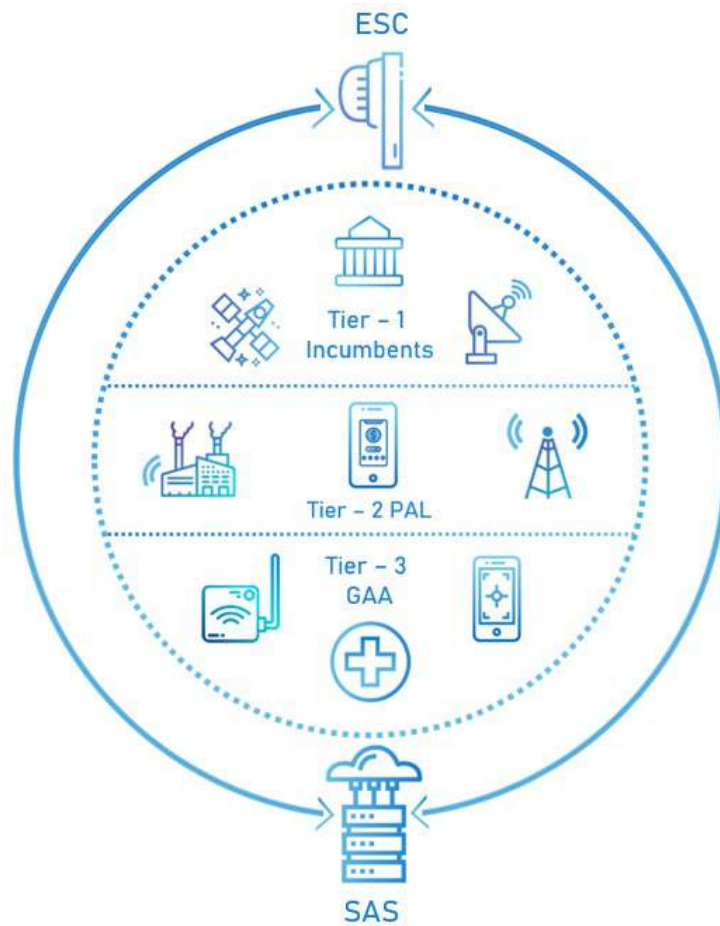
### 2.1. Citizens Broadband Radio Service

On recommendation of President's Council of Advisors on Science and Technology (PCAST) board advisory to lead innovation in wireless technology services, the U.S. government directed National Telecommunications and Information Administration (NTIA) to collaborate with Federal Communications Commission (FCC) to find possibilities for the commercial use of additional spectrum, which held by Federal and non-federal users. In 2015, the FCC adopted rules for commercial use of 3.5 GHz frequency spectrum in the range from 3550 to 3700 MHz for shared wireless access. The commission established Citizens Broadband Radio Service (CBRS) and created three-tier framework to accommodate shared frequency band for the commercial use. CBRS network devices use the LTE standards for radio access and core network as regular LTE devices. The main difference between traditional LTE and CBRS LTE operation is the spectrum access method that devices use. In normal LTE operation, wireless service operators buy license LTE spectrum and utilize the spectrum for an access. In CBRS LTE, a central spectrum assignment mechanism used to manage the spectrum utilization and access. This new central spectrum assignment mechanism known as a Spectrum Access System (SAS)

### 2.2. A novel three-tier paradigm, an industry first with CBRS



**Figure 1 CBRS 3.5 GHz Frequency Band**



**Figure 2 3-Tier CBRS Structure**

A three tier-sharing framework, SAS enforced prioritization at all times.

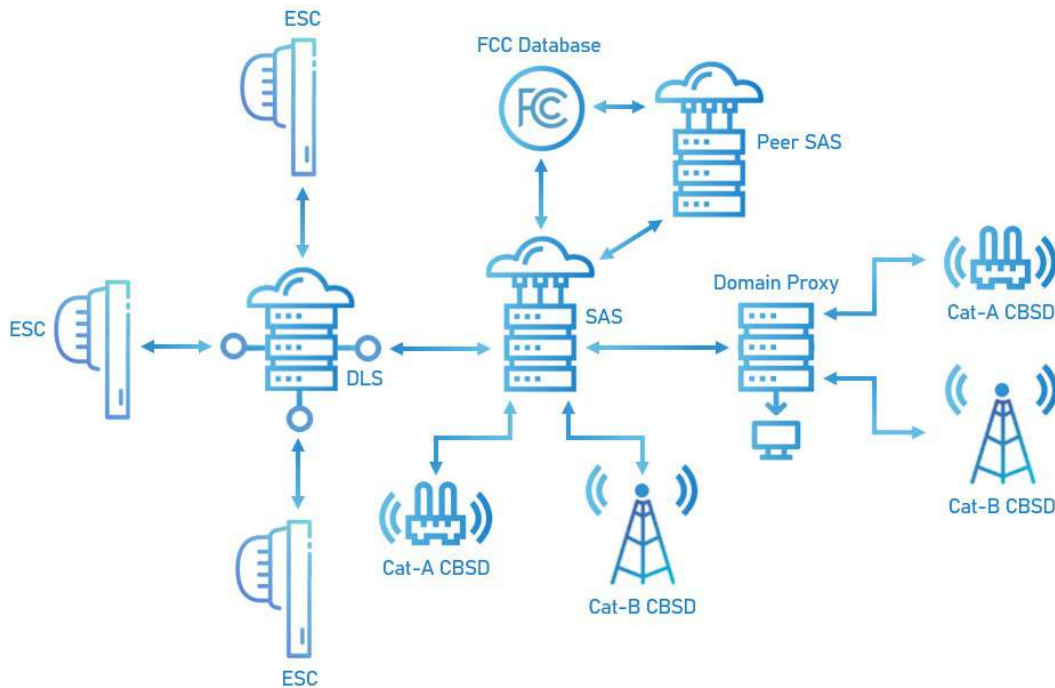
**Incumbents:** US Military Radar/DoD, Fixed Satellite Station (FSS), Grandfathered Wireless Broadband Licensees (GWBLs) Highest priority access over PAL and GAA users.

**PAL:** Priority Access License, licensed tier. Licenses awarded via FCC Auction 105 for 3.5 GHz. Priority access to 70 MHz. Must register with SAS. Interference protected from GAA users.

**GAA:** General Authorized Access, unlicensed tier. Minimum of 80 MHz with an ability to float up to 150 MHz if PAL licenses are unoccupied. Must register with SAS. No interference protection. Opportunistic use, null priority rights.



### 2.3. Spectrum Access System



**Figure 3 Spectrum Sharing Ecosystem**

ESC: Environmental Sensing Capability  
DLS: Decision Logic System  
SAS: Spectrum Access System  
CBSD: Citizens Broadband Radio Service Device

Spectrum Access System (SAS) is an entity authorized by the Commission to operate in accordance with the rules of part 96. It manages the assignment and distribution of spectrum to the users based on permissible frequencies at their location. Depending upon the user location, SAS controls and permits maximum transmission power level to protect higher tiers incumbent users and PAL users. All category CBSDs must register, authenticate identification and report its location with the SAS. An individual CBSD may directly communicate with SAS or CBSDs composed of a network may employ a system (Domain Proxy) for aggregating and communicating all required information exchange between CBSDs and SAS. On exchange of all necessary CBSD's information with the SAS, the SAS verifies requestor CBSD's FCC identifier with a FCC database and verifies CBSD's identity prior to authorizing the state where CBSD can start providing services. SAS is capable to determine available channels for users at any given geographic location and on the request of Spectrum Inquiry from CBSDs; SAS shall respond back to CBSDs with available spectrum for the requesting CBSD's geo location so CBSD can proceed with Grant request to SAS. CBSD would look for a channel grant from SAS by providing its maximum EIRP capability per MHz and on conforming to CBSD's EIRP capability with FCC database; SAS shall approve or reject Grant for requesting CBSD. For any new requesting CBSD, on successful Grant response, SAS would put CBSD in IAP pending until next Coordinated Periodic Activates among SASs (CPAS) window before giving Authorization to such CBSD. CPAS runs through every night from 12am PST to 3am PST. Post CPAS window; granted CBSD would be authorized to transmit by SAS followed by conditional exchange of heartbeats between CBSD and SAS at every 200 seconds time interval.

## **2.4. Environmental Sensing Capability**

Environmental Sensing Capabilities (ESCs) sensors used to detect operation of Federal Incumbent in 3550 to 3700 MHz frequency band and it provides incumbent information back to SAS upon detection of any activity on any of the channels in the band. On receiving Incumbent detection information from ESC, SAS shall direct CBSDs for channel change or cease CBSDs transmission by suspending active grants of CBSD in order to protect federal incumbent's operation. On active Grant suspension by SAS, CBSD shall relinquish that grant and request new Grant for alternate channel recommended by SAS for immediate authorization. If no alternate channel available then SAS shall cease CBSD transmission by suspending its last active grant. On suspended grant, CBSD may continuously send heartbeats to SAS for suspended grant and resume operation when it receives success code (Code "0") from SAS in any of the heartbeat responses from SAS. In Release-1 of CBRS specifications, ESCs only monitor activity of Federal incumbent in 3550 to 3650 MHz of spectrum. ESC is a non-federal entity commonly operated by approved SAS administrators. Due to ESC's importance of detecting federal incumbents, SAS also protects ESCs against harmful interference which may be caused by surrounding CBSDs.

## **2.5. Incumbents**

Incumbents are Tier-1 users of the CBRS frequency spectrum with highest priority over PAL and GAA users. There are Federal and Non-Federal incumbents. Federal incumbents are U.S. Navy's Shipborne Radar System that operates on any frequencies between 3550 to 3650 MHz. Federal incumbents also include some inland radar operation sites operating on any frequencies between 3500 to 3650 MHz. Non-Federal incumbents include Fixed Satellite Stations (FSSs) and Part-90 Grandfathered status Wireless Internet Service Providers. They are called as Grandfathered Wireless Broadband Licensees (GWBLs). FSSs are receive only earth stations operating on frequencies between 3600 to 3700 MHz. There are a few Out-of-Band (OOB) Registered for protection earth stations, i.e. TT&C (Telemetry, Tracking and Command) FSSs operate on frequencies between 3700 to 4200 MHz. GWBLs operate only on upper 50 MHz of CBRS spectrum, i.e. 3650 to 3700 MHz. These are all Tier-1 incumbents in CBRS shared model eligible for interference protection by SAS. Spectrum controller system retains information about these Federal as well as Non-Federal incumbents and their protection and exclusion zones in accordance with the rules of part 96. Depending upon CBSD's geo location and relation between non-federal incumbents (location of FSS and GWBL's active transmitter), SAS shall determine if the deployed CBSD location is an exclusion or protection zone for that CBSD to protect that non-federal incumbent from harmful interference. SAS also retains information about list of Exclusion Zones (EXZs) published by NTIA and shall not allow CBSD operation within 40kms radius of defined locations.

## **2.6. Protection of Federal Incumbents**

Federal incumbents are-protected by SAS through the method of Dynamic Protection Area (DPA) activation. There are two types of DPAs. E-DPA and P-DPA. E-DPA is an ESC-DPA and an always-on DPA, monitored by ESC network. NTIA has provided the list of E-DPAs along the coastal boundary in the east and west coast of U.S. geography. Other than coastal areas, the additional list of E-DPAs provided by NTIA also covering the areas of Alaska, Hawaii, Puerto Rico and Guam. When federal Navy's radar system turn on its operation then SAS administrator's ESC sensor detects the active channel of radar system and the DPA zone where Navy ship operates. Upon detection, ESC provides that DPA ID and channel information to SAS. So within 300 seconds SAS shall re-assign CBSDs those are operating on the same channel as Navy's radar system and are under the protection zone of that DPA if there is any alternate channel available in that geographical area or may even cease those CBSDs' operation if no alternate channel available in order to protect federal incumbent from harmful interference. In order to protect inland federal incumbents, P-DPAs (Portal-DPAs) are used. There are few P-DPAs published by NTIA for

federal's ground based radar test sites operate on frequencies in the range of 3500 to 3650 MHz. Also known as "Informing incumbents" where SAS gets testing schedule of ground based radar system 24 hours in advance. In order to protect in land federal incumbent, SAS shall re-assign CBSDs operating on the same channel as ground radar system's channel and located under that P-DPA zone.

## **2.7. Protection of Non-Federal Incumbents**

Non-Federal incumbents such as FSSs and Part-90 GWBLs protected by SAS from harmful interference caused by low tiered users through the method of aggregate interference calculation as per rules defined in part 96. Aggregate interference calculations consider CBSDs within 40 to 150 kms radius depending on the type of protected entities and type of CBSD.

## **2.8. Commercial SAS Trial Observations in New York City, NY**

Charter deployed a cluster of 21 FCC certified Category-B outdoor CBSDs in New York City, NY area. All CBSDs are configured with single 20 MHz LTE channel transmitting at max power per Category-B CBSD limit defined in part 96 under general radio requirements. New York City, NY county area is within 150 kms radius of one of the registered for protection "Fixed Satellite Station" (FSS) located in The Bronx, NY. There is an unexpired GWBL's transmitter located within 40 kms radius of this FSS. Hence, this relation between FSS and GWBL makes the area of 150 kms radius of the FSS an "Exclusion Zone". Therefore, upper 50 MHz of spectrum found to be unavailable for CBRS services by SAS until the expiration of active GWBL's license. So CBRS services can only be operational in the spectrum between 3550 to 3650 MHz in the New York City county geographical area. As per NTIA's defined protection zones, the deployed geographical area covered by five E-DPAs and four P-DPAs. CBSD grant suspension observed by any of these five E-DPA activations whenever there is any incumbent activity detected by ESC network in these E-DPA zones. The only channel suspended by SAS where Navy's radar system is operating on. Due to unavailability of alternate channel re-assignment feature on SAS now, CBSD could not get instantaneous new grant authorization from SAS. CBSD stayed in grant suspended state while Navy's radar system was in operation and resumed its services back when SAS gets no further information from ESC about incumbent activity.

## **2.9. Conclusion**

CBRS is technology agnostic so it's not only good for LTE services but also one of the best option to deploy 5G services today in mid-band range of spectrum because of possibility to have wider channel bandwidth up to 100 MHz in 5G as per FR1 specs of 3GPP. Use of CBRS will provide platform for high performance deployments of diverse use-cases. It can be as simple as point to point or point to multipoint deployments for a small networks to large complex networks. These use-cases can be fixed, mobility and even advanced standalone networks. Other than mobile and cable service providers, many other market verticals such as Medical, Industrial, Agricultural, Retail, Oil & Gas, Energy, Power utility, Transportations, Airport and Educational institutions can utilize CBRS spectrum and make best use of it for their private network, Internet of Things (IoT), Security and surveillance and much more.

### 3. Implications to the deployment strategies


This section covers the details of how Charter has leveraged it's all possible physical assets such as Cable strand, Towerstream buildings and valued SMB locations for strategic deployment of its wireless mobility network in CBRS across 5 different markets. The deployed networks have been fully utilized for Charter's extensive in depth testing of MVNO data offload use case.

#### 3.1. Attached Mount Deployment (Rooftop)

47 dBm/10 MHz allowable power limit applies to this deployment type and it can support larger antennas e.g. 64T64R Massive-MIMO antennas, one CBSD can have multiple sectors, typically get mounted on the rooftops and connect to Charter's DOCSIS serving the building.

**Application:** Attach mount unit can be installed more strategically than strand mount or SMB to clear obstructions or point to a targeted hotspot with required down-tilts without limitations like the other two types (Strand, SMB). Attached mount will be more effective when advanced features like Antenna Beamforming, Sector Virtualization, and Dynamic Load Balancing etc. are tested. This type of deployment provides larger coverage and capacity than strand and SMB scenarios.

**Cost:** Attach mount deployment and radio costs are most expensive than other three types of deployment. High cost for this type of deployment due to site surveys, A&E (structural analysis, construction drawings, LPC, transit or any other permit required) and entering into lease agreement with landlord. Site installation and commissioning, city inspection and close-out.



**Figure 4 Attached Mount (Rooftop) Outdoor Deployment**

#### 3.2. Strand Mount Deployment

Strand mount is the most cost effective solution for Charter where aerial cable strand lines exist. The size and power are lower than Attached mount CBSDs, but their biggest bottleneck is power consumption from

the HFC plant power supply. Charter's mandating vendors to stay under 100W, some vendors have come up with strategies like powering down the CBRS amp until traffic picks up to keep plant power consumption low.

**Application:** Another bottleneck of strand unit is the mounting orientation – It has to be always along the strand and thus hotspot-targeted deployment in this case might be challenging. To mitigate this, Charter has requested Quasi-Omni strand design that has dual sectors with two sets of antenna covering NE and SW directions. Their height is always 18ft and typically, comes with 2x2 MIMO capability.

**Cost:** Utilizing existing assets to mount, connect to, maintain and operate make these the most cost efficient and expedient deployment type for outdoor Cat-B CBSDs.




**Figure 5 Strand Mount Outdoor Deployment**

### **3.3. Small Medium Business (SMB) Deployment**

Indoor low power Category-A devices have been installed as part of Charter's CBRS wireless trial on Charter business customer's premises. Mainly used to provide pedestrian coverage. This deployment type can provide blanket coverage when deployed every 400 - 500 ft. A Strand or Attach mount usually compliments SMB coverage by serving as an umbrella cell and fill in for any coverage holes. Charter has requested vendors for Tri-Star config which adds a third sector to cover indoor and two sectors pointing outside the stores from behind a glass window. Charter makes use of this deployment type where applicable to form a uniform layer of CBRS coverage targeting AOI's and have run trials confirming seamless connectivity and performance between the outdoor coverage and the indoor coverage.

**Inside Out:** Charter's initiative to provide outdoor / pedestrian street coverage utilizing business customers' locations tested to serve as a contingent layer of indoor Category-A devices under outdoor high power Category-B umbrella cell. Below is a snapshot of three different locations where Indoor units installed with

different heights and attenuation after the glass within 100ft was observed to be least at 18ft. This resulted in a larger footprint and relatively sustained coverage of the three scenarios.



**Figure 6 Indoor CBSD Deployment in SMB**



**Figure 7 Inside-Out Deployment in Spectrum Stores**

**Table 1 Inside-out Testing**

Inside-Out Testing	Manhattan Spectrum Store	Astoria Spectrum Store	7th Floor Charter Office
Coverage Radius	1000 ft.	485 ft.	495 ft.
Throughput Outside	67 Mbps / 19 Avg.	74 Mbps / 17 Avg.	49 Mbps / 9 Avg.
Throughput Inside	77 Mbps	74 Mbps	72 Mbps
RSRP within 100ft.	-78 dBm	-95 dBm	-102 dBm
Height	18 ft.	8 ft.	70 ft.
Clutter Type	Dense Urban	Urban	Dense Urban
Glass Type	Standard	Double glass panel	Standard

**Conclusion:** Inside-out strategy requires optimal height (around 15 ft.) non-reflective and non-metal coated glass type and strategic placement of CBSD behind the glass in SMB location.



One of the most frequent issue noticed on SMB type indoor small cell is TDD time synchronization using external independent GPS antenna due to unreliable GPS signal with low SNR level in dense urban and urban locations in New York City and Los Angeles. It has been observed that small cells keep going into a GPS holdover state and out-of-sync state followed by out of service very frequently. It is unrealistic to use IEEE 1588PTP solution with independent grand master clock on each and every SMB locations. SMB small cells are wired using DOCSIS3.1 indoor CM for backhaul connectivity and hence it will be essential to have DOCSIS Timing Protocol (DTP) on such locations to mitigate timing sync related issues. DTP has been described in detail in section 8.1 of this document.

**Outside-In Tests:** The nature of the 3.5 GHz wave makes it less reliable when coverage target area is indoor and CBRS cell deployed outside. We conducted tests with a few types of inbuilding locations with an outdoor CBRS serving cell on Strand. Our tests include, location with Brick Wall and no windows in corridor, location with windows but obstructed for outside signal and location with windows and no obstructions to outside signal.



**Figure 8 Outside-In Test Locations**

**Table 2 Outside-In Testing**

Outside-In Testing	High School (Inside)	Grocery Store (Inside)	Restaurant (Inside)
RSRP (dBm)	-112 dBm (best), -135 dBm (worst)	-118 dBm to -135 dBm (outside -99 dBm)	-105 dBm (outside -100 dBm)
Throughput (Mbps)	9 Mbps (avg), 0 Mbps (min)	8 Mbps (outside 39 Mbps)	30 Mbps (outside 39 Mbps)
Coverage Limiting Factor	Uplink timeout	20-30 dBm Loss of RSRP	Glass type
Structure Type	Brick walls	Glass window + wall	Full glass windows
Clutter	Residential/Urban	Commercial/Urban	Commercial/Urban
Will Outside-In strategy work?	No	Weak	Yes (Rare case)

**Conclusion:** Providing the coverage inside the building from outdoor CBSDs will not work effectively. If Area of Interest (AOI) is inbuilding then it has to be targeted within the building.

## 4. Performance Characteristics of CBRS LTE Network

CBRS is a beneficial platform to bridge the gap between low spectrum carriers and mmWave challenges. Below are some of the characteristics Charter has observed in their CBRS trials.

### 4.1. Characteristics of CBRS Wave

3.5 GHz electromagnetic waves are subject to increased propagation losses than say the signals of AWS 1700 MHz or GSM at 800 MHz in a cellular network. The CBRS wave is 8.5 cm in wavelength, compared to 17 cm (AWS) and 37.5 cm (GSM) at the above frequencies. Thus, a CBRS cell gets- 2 times more propagation loss from 1.7 GHz carrier and 4.4 times higher propagation loss from 800 MHz frequencies. Extrapolating these values, if we suppose 3.5 GHz Category-B CBSD's Cell Radius in a dense urban clutter is 1/4<sup>th</sup> mile (1320 ft.) would yield cell radius of AWS radio at about 1/2 mile (2650 ft.) and about 1.1 miles (5808 ft.) for GSM keeping same EIRP.

Antenna size advantage: CBRS has a cutting-edge advantage over low band carriers making it more favorable to rollout. Due to smaller wavelength, antenna size is typically 12" – 15" long and weighs 2 lbs. This has opened-up deployment possibilities /strategic venues that were not an option in lower band rollouts from an acquisition standpoint. Image below shows CBRS 3.5 GHz antenna to the left and 600 MHz antenna on the right on one of Charter's trial sites. CBRS antenna is 14.3" long, wights 2 lbs compared to 600 MHz antenna's 8 ft length and more than 100 lbs weight.

Frequency Range	MHz	3300 - 3800MHz
Polarisation	Degree	+/-45° Slant Linear
Gain	dBi	12.5
Azimuth Beamwidth	Degree	65°
Azimuth Beam Squint	Degree<	3°
Elevation Beamwidth	Degree	22°
Electrical Downtilt	Degree	T0°
Electrical Downtilt Deviation	Degree<	1°
Impedance	Ohms	50
VSWR	<	1.5
Return Loss	dB>	14
Isolation	dB>	25
Front to Back Ratio: Total Power +/-30°	dB>	28
Upper Sidelobe Suppression, Peak to 20°	dB>	18
Cross-Polar Discrimination	dB>	16
Maximum Effective Power Per Port	W	50



**Figure 9 CBRS 3.5GHz Antenna Properties**

**Decreased Obstruction Resistance:** CBRS wave is seen to incur abrupt attenuation when obstructed compared to AWS and GSM that could refract around obstructions, form a shadow and resume coverage in good RSRP range. CBRS however is observed to be very prone to obstructions. This can be seen in drive data when the cell is obstructed the coverage takes an abrupt hit i.e. 20-30 dBm attenuation.

**3.5 GHz coverage reliability:** CBRS wave is affected to a higher degree relatively when deployed in high foliage areas, hilly terrain obstructing antenna beam, downtowns or vertically layered dense urban clutter that obstructs the beam.






**Figure 10 Dense Urban Morphology**

## 4.2. mmWave Comparison

If an operator doesn't already own spectrum, then there are a few options which can be considered. The next auction which will be happening is the C-BAND auction. This spectrum will have very similar characteristics to CBRS and should be viewed in a similar manner. The other option is to look at mmWave. mmWave has two distinct advantages. The first advantage is large blocks of bandwidth which can be utilized versus sharing 150 MHz among many parties. The second advantage is the ability to radiate at higher powers due to the restrictions the FCC has placed on CBRS. Both of these are strong advantages for mmWave; however, mmWave has a major issue. It does not propagate very far, and it is easily obstructed by foliage, buildings, and windows. To show the differences, a theoretical comparison was performed between 5G CBRS and 5G 28 GHz mmWave. Table 3 below shows the baseline assumptions for the comparison and Figure 11 shows the results.

**Table 3 5G CBRS vs 5G 28 GHz mmWave Assumptions**

	mmWave gNB	Sub-6 gNB	mmWave UE	Sub-6 UE
Combined transmission power (dBm)	28	34	16	23
Antenna gain (dBi)	23	23	11	0
Total EIRP (dBm)	51	57	27	23
BW (MHz)	800	100	800	100
Frequency (GHz)	28	3.5	28	3.5
Height (m)	40	40	1.5	1.5
DL/UL max layers	2/2	8/4	2/2	2/2
Propagation Model	Rma Rural NLOS	Rma Rural NLOS	Rma Rural NLOS	Rma Rural NLOS



**Figure 11 5G CBRS vs 5G 28 GHz mmWave Results**

**Key takeaways from the results:** DL throughput is not significantly higher for the 28 GHz mmWave transmission versus the 5G CBRS product even though the bandwidth is 8X. 5G CBRS gets significantly more coverage (1.7 miles versus 0.5 mile); however, UL connectivity is very low just pass 0.5 mile. UL capacity below 0.2 mile is significantly greater in mmWave.

### 4.3. 5G CBRS vs 4G CBRS

Outside of the hype, does it make sense to deploy a 5G CBRS network versus a 4G CBRS network? There are many areas which need to be explored before making this decision. For the purposes of this paper, a short review of a few of the technical aspects will be performed.

From a theoretical comparison, spectral efficiency for 5G is slightly better due to a leaner carrier design which corresponds to roughly 15% improvement. Very similar improvement for 5G is also seen with Pathloss calculations.

General comparison between LTE and 5G –

**Table 4 LTE vs NR Comparison**

	LTE	NR
Frequency Range	Sub-6 GHz	Sub-6 GHz and mmWave
Maximum CC BW	20 MHz	100 MHz or 400 MHz depending on frequency band
Subcarrier Spacing	Fixed 15 kHz	Scalable $2^{\mu} 15$ kHz
Waveform	DL: CP-OFDM UL: DFT-s-OFDM	DL: CP-OFDM UL: CP-OFDM/DFT-s-OFDM
TTI	14 OFDM symbols in fixed 1ms	14 OFDM symbols with scalable $1/2^{\mu}$ ms
Channel Coding	Data: Turbo coding	Data: LDPC Control: Polar code

	Control: Tail biting convolutional code	
Initial Access	Broadcast	Unicast
MIMO	8 layers codebook and non-codebook precoding	8 layers non-codebook precoding
Reference Signal	CRS, DMRS, CSI-RS, SRS	DMRS, CSI-RS, PTRS, SRS
Duplexing	FDD, Static TDD	FDD, Dynamic TDD
HARQ	Synchronous/Asynchronous with 8 processes	Asynchronous with 16 processes and CBG retransmission

#### 4.4. CBRS Coverage Reliability from Field Tests

Charter conducted tests to observe CBRS behavior and coverage reliability in different clutter types. It has been noted that CBRS is more prone to obstructions and cell radius is affected when path is obstructive. Phenomenon like refraction and scattering do not aid CBRS wave in an obstructed path.



**Figure 12 Field Drive Test Plot**

Site data shown in the Figure 12 where street to the north is obstructed. It has been observed the RSRP cannot sustain obstruction and degrades to sub -120 dBm zone very quickly right after (138 ft). Unobstructed path to the west is 736 ft. SINR also shows similar trend to RSRP degradation. This is typical behavior of CBRS 3.5 GHz wave when obstructed and reduces effective cell radius in non-LOS conditions.




#### 4.5. Factor affecting throughput in CBRS network – Field Assessment

Charter tested various factors to observe effects on throughputs. Below is the analysis of data gathered. From the field we observed DL throughput has high dependency on RSRP. As CBRS RSRP fluctuates the affect is translated directly to the streaming UE. As seen from test results, depending on morphology of the clutter, below a certain RSRP the SINR takes a hit and throughputs plunge to single digits.



**Figure 13 Field Single Site Performance Stats**


In the left graph: As highlighted, when RSRP (blue) drops below -115 dBm, the DL Throughput is seen taking an abrupt hit (gray) to single digits. It rises again as RSRP goes up and again below certain RSRP the throughput nose dives. The second highlighted region from left to right, shows obstruction in coverage as RSRP taking a steep gradient downward, the throughput degrades but doesn't drop down to minimum until RSRP crosses neg 115 dBm. For this morphology, neg 115 dBm seems to be the point below which decent throughputs cannot sustain. In the right graph: In the second site region the throughput shows same trend as first region, as shown in highlighted area, as RSRP (blue) dips below neg 100 dBm, the SINR drops and Throughput take a steep slope downward to single digits (gray). As soon as RSRP exceeds -100 dBm, both SINR and throughput rise showing direct correlation. The second highlighted area shows similar trend when RSRP drops.



**Figure 14 Field Cluster DL Performance Stats**

In another area cluster drives were observed to show instances (left highlighted portion above) where RSRP stayed between -110 dBm to -112 dBm and throughputs were not severely affected even when SINR dropped to single digits. The right highlighted portion shows an abrupt drop in DL throughput as RSRP plunges but it regains and trends alongside RSRP rising trend.

Also this area was deployed with a different vendor equipment than the previous area so the chipset or equipment's receive sensitivity or processing power could be the factors differentiating from other area. In addition to this, this clutter morphology is denser with more shadowing, reflections, refractions and beam scattering phenomenon(s) coming into play.



**Figure 15 Field Cluster UL Performance Stats**

In this graph UL throughput data is shown, collected during a field trial. It can be seen UL is more susceptible to degrade as RF conditions deteriorate. When DL RSRP dips below -100 dBm, the UL throughput lowers to under 1 Mbps. Although DL RSRP is uncorrelated to UL but the path loss scenario's are the same in both directions most of the time, as highlighted in the graph above show.

#### 4.6. LTE Performance KPIs Observation from Field Trials

Performance KPIs for Strand network when tested yielded following –

Up to 14 Simultaneous UEs connected in dedicated mode, 14 GB Tonnage carried with 1% Drop Rate and 99% Establishment Rate and X2 HO Success rate of 95%

**Table 5 Cluster Performance KPIs**

Strand CBRs Cluster KPIs	Max UEs Connected	Drop Rate (%)	ERAB Success Rate (%)	RRC Success Rate (%)	S1 Success Rate (%)	Tonnage (GB)
	14	1	99	99	100	14

Attached mount cluster field tests showed following –

#### Cluster KPIs

Success Rates – RAB, RRC, HO

ERAB Success Rate (%)	RRC Success Rate (%)	S1 Success Rate (%)	Inter-ENB HO Success Rate (%)	S1 HO Success Rate (%)	X2 HO Success Rate (%)
99.09	99.75	100.00	100.00	100.00	100.00

Establishment

ConnEstablishAttSum	ConnEstablishSuccSum	InitialAttemptedToSetupSum	AdditionalSuccessfullyEstablishedSum	InitialSuccessfullyEstablishedSum
963	956	929	0	924

HO Statistics

UeAssoc S1ConnEstablishAtt	UeAssoc S1ConnEstablishSucc	InterEnb OutPrepAtt	InterEnb OutSuccSum	S1 InterEnb OutPrepAtt	S1 InterEnb OutSuccSum	X2 InterEnb OutPrepAtt	X2 InterEnb OutSuccSum
956	956	35	35	1	1	34	34

**Figure 16 Attached Mount Cluster KPIs**


In a single session 5.9 GB network tonnage was handled with 43 HO attempts in between 3 sites, 8 Sectors and 16 Carriers. Up to 14 UEs, 5 simultaneously attached to the network, peak data rates of

103 Mbps observed. 963 connection establishment attempts, ERAB setup rate, RRC Success rate, S1 & X2 HO Success rate and Inter-eNB success rates were all greater than 97%.

RSRP	Percentage of Drive (%)	DL Thpt Avg (Mbps)
> -90 dBm	11	41
> -100 dBm	17	35
> -110 dBm	37	25
> -120 dBm	21	16
> -130 dBm	14	5

**Figure 17 Cluster Drive RSRP and Throughput**

RSRP is the key driver of performance. Below graph is a plot of DL throughput per dBm and the steepest gradient is at -104 dBm. We've observed depending on clutter type and morphology the throughput is more affected as RSRP degrade from -104 dBm and drops to single digits as it nears -120 dBm as shown in 'Factor Affecting Throughput in a CBRS Network – A Field Assessment' section.



**Figure 18 RSRP to Throughput Function**

#### 4.7. Capacity Testing

Charter conducted a field trial in Tampa, FL to test CBRS load bearing capability and effect on cell performance throughput. Field trial done at three distinct locations from the cell based on RSRP values referred as Cell Near, Cell Mid and Cell Edge.






Figure 19 Load Test Setup on Field




Figure 20 Load Test DL Performance Stats

In the graph above, when one UE attached at Cell Near location, the Downlink throughput was 60 Mbps. When second UE was fired up the Downlink throughput dropped to 35 Mbps Avg. Third one joined, and the throughput reduced to 27 Mbps and so on (blue). At the Cell Near location, up to 36 UEs simultaneously streamed above 2Mbps. At the Cell Mid location (orange), throughputs started at 33 Mbps and up to 27 UEs could stream above 2Mbps simultaneously. Lastly, at the Cell Edge (gray) location, 20 Simultaneous UEs could stream at and above 2 Mbps. Testing of another vendor showed slightly different results, consolidating all in the table below shows Cell Near of 34 and Cell Edge of 13 UEs maintaining 2 Mbps at all times.



**Table 6 DL Load Testing**

DL Throughput	# UEs at Cell Near	# UEs at Cell Edge
> 20 Mbps	4	0
15 – 20 Mbps	2	0
10 – 14 Mbps	4	0
5 – 9 Mbps	7	4
2 – 4 Mbps	<b>34</b>	13
0 – 1 Mbps	0	<b>33</b>

Majority of UEs stayed within 2 – 4 Mbps at Cell Near.

Majority of UEs stayed under 1 Mbps at Cell Edge

In addition, from observation in the Field trials that as UEs are loaded up, the 2Mbps boundary threshold (DL) reduced virtually. Our test results shown in the table below:

**Table 7 Load Effect on Cell Radius**

Loading effect on Geographical Radius of 2 Mbps DL boundary		
# of Simultaneous UEs	2 Mbps Radius (m)	% Loss of Radius
1	342	
5	277	19%
15	205	40%
30	64	81%

#### 4.8. Busy Hour Traffic Analysis

Busy hour changes dynamic of the cell, grade and quality of service. Below is a theoretical analysis of how busy hour affects end user throughputs. Input parameters in this calculation have been taken from our capacity testing field trial.

We assumed average throughput per user is 50 Mbps DL / 10 Mbps UL in non-busy hour. Using capacity trial results, we get average of 34 Cell Near and 13 Cell Edge, giving an average of 23.5 UEs Cell Mid for 20 MHz. At 40 MHz we took 1.5 times 23.5 UEs = 35 UEs per cell. Assuming 80% of users at busy hour, this comes to 28 UEs per cell in BH. Combining rates for different QAMs, throughput per user in BH comes out to be 5.8 Mbps. From 50 Mbps to 5.8 Mbps is the effect of busy hour on a CBRS cell. This analysis is theoretical.

**Table 8 Assumptions for Busyhour (BH) effect on Throughput**

Per CBRS Cell 40 MHz BH Analysis	DL	UL	Unit	Remark
Avg. Throughput per UE (Product)	50	10	Mbps	20% UL based on FC2
# of UEs per Cell	35		Avg	Per testing in 20 MHz: CN 35 + CE 12 = 23.5 Avg UEs * 1.5 for 40 MHz = 35.2 UEs
BH Users	28		UEs	80% of estimated total UEs
Peak Throughput BH (256 QAM)	300		Mbps	
Peak Throughput BH (64 QAM)	220	35	Mbps	
Peak Throughput BH (16 QAM)	150	23	Mbps	
Peak Throughput BH (QPSK)	75	12	Mbps	

Peak Throughput BH (Consolidated)	164	17	Mbps	(Based on drive data: 30% for each DL Modulation except 10% for 256 QAM)
Throughput per User (BH)	5.8	0.6	Mbps	
Total Tonnage per User (BH)	2.6	0.3	GB	
Total BH Tonnage per Sector	73.6	7.5	GB	

**Table 9 Busyhour (BH) Traffic Analysis - Strand Mount**

<b>BH Traffic Analysis – Strand</b>	<b>DL</b>	<b>UL</b>	<b>Unit</b>
Avg. Throughput per UE (Product)	50	10	Mbps
# of UEs per Cell	35		Avg
BH Users	28		UEs
Peak Throughput BH (Consolidated)	164	17	Mbps
Throughput per User (BH)	5.8	0.6	Mbps
Total BH Tonnage per Sector	73.6	7.5	GB
Total BH Tonnage per Strand Site	147.2	15.0	GB

**Table 10 Busyhour (BH) Traffic Analysis - Attached Mount**

<b>BH Traffic Analysis – Attached</b>	<b>DL</b>	<b>UL</b>	<b>Unit</b>
Avg. Throughput per UE (Product)	50	10	Mbps
# of UEs per Cell	35		Avg
BH Users	28		UEs
Peak Throughput BH (Consolidated)	164	17	Mbps
Throughput per User (BH)	5.8	0.6	Mbps
Total BH Tonnage per Sector	73.6	7.5	GB
Total BH Tonnage per Attached Site	220.7	22.5	GB

**Table 11 Busyhour (BH) Traffic Analysis - SMB**

<b>BH Traffic Analysis – Attached</b>	<b>DL</b>	<b>UL</b>	<b>Unit</b>
Avg. Throughput per UE (Product)	50	10	Mbps
# of UEs per Cell	17		Avg
BH Users	14		UEs
Peak Throughput BH (Consolidated)	164	17	Mbps
Throughput per User (BH)	12.0	1.2	Mbps
Total BH Tonnage per Sector	73.6	7.5	GB
Total BH Tonnage per SMB Site	220.7	22.5	GB

Average DL User throughput drops from 50 Mbps to 5.8 Mbps as traffic spike hits 80% of cell capacity. Cell overlap along with advanced load and spectrum sharing and congestion mitigation techniques can be employed to reduce this impact.

## 4.9. Conclusion

CBRS is more suited for hot-spot targeted and pocketed coverage applications with redundancy to fall back to alternate technology when needed like Charter's MVNO in DSDS configuration with the ability to transition to MVNO network to mitigate customer impact. In the event of blanket / cluster wide coverage requirement the ISDs will have to be no more than 1000 – 1400 ft. apart.

CBRS LOS vs non-LOS coverage, the loss beam incurred when subjected to obstruction due to its high propagation loss characteristics, makes CBRS radius shrink to a greater extent than low band spectrum carriers.

## 5. CBRS Radio Devices and RF Design of CBRS Network

### 5.1. CBRS Radio Devices

In CBRS terminology, an LTE eNodeB or base station is called Citizens Broadband Radio Service Device (CBSD). As per Code of Federal Regulation, Part 96, there are two categories of CBSDs defined, Category-A and Category-B devices. All CBSDs must register with and authorized by SAS prior to their initial service transmission. They must operate at or below the maximum power level authorized by SAS and must be in compliance with their FCC equipment authorization. For Category-A CBSD, the maximum EIRP and maximum PSD limits are 30dBm/10MHz and 20dBm/MHz respectively and generally deployed indoors. For Category-B CBSD, the maximum EIRP and maximum PSD limits are 47dBm/10MHz and 37dBm/MHz respectively and deployed outdoors.

### 5.2. CPE for Fixed Wireless Access

Selection of optimal Customer Premise Equipment (CPE) for Fixed Wireless Access (FWA) application is very crucial. High power, high antenna gain CPE device helps in network design for better uplink user performance perspective. Today, there are many OnGo certified, part 96 compliant CPEs available in the market to choose from for various different form factors for Fixed Wireless network application in CBRS. Charter has done FWA rural broadband trial in North Carolina with a couple vendor's CPE devices with capability of 1x4 and 2x4 MIMO configurations and high power up to 26 dBm and 15 dBi high gain antenna for better user experience. Their 3GPP device category information is very important and used to allow the eNodeB to provide compatible user services more efficiently. In other words, the user equipment category defines the overall performance and the capabilities of the UE. There are different UE categories that have a wide range of features supported for enhanced end user performance experience. Some of them are only capable of supporting SISO; some of them support 2x2 and 4x4 MIMO. The UE category defines a combined uplink and downlink capability as specified in 3GPP [TS36.306](#)

**Table 12 CPE Category**

User Equipment Category	Downlink (Mbits/s)	Max # of DL MIMO Layers	Uplink (Mbits/s)	3GPP Release
1	10.3	1	5.2	Rel. 8
2	51	2	25.5	
3	102	2	51	
4	150.8	2	51	
5	299.6	4	75.4	

User Equipment Category	Downlink (Mbits/s)	Max # of DL MIMO Layers	Uplink (Mbits/s)	3GPP Release
6	301.5	2 or 4	51	Rel. 10
7	301.5	2 or 4	102	
9	452.2	2 or 4	51	Rel. 11
10	452.2	2 or 4	102	
11	603	2 or 4	51	
12	603	2 or 4	102	

### 5.3. RF Design

Various different strategic CBRS trials have been performed over the last couple of years. Extensive Continuous Wave (CW) tests performed in different morphological geography to understand 3.5 GHz signal propagation characteristics in all different areas like dense urban, urban, sub-urban and rural. Results have been studied thoroughly, data utilized for RF design model tuning to create our own optimal propagation model for each different morphological area. For an accurate and realistic RF design, the recommendation is to use 3D modeling with accurate and high definition 3D geo data for a better result. Foliage data consideration into RF design would give result that is more accurate. For the best RF design and propagation model tuning, adopt an advanced method to collect CW data with better equipment. The advanced method such as dead reckoning giving location information that is more accurate. This will be crucial for markets such as NYC, LA and Dallas or any downtown areas with high-rise buildings.

### 5.4. Conclusion

Plenty of RF drive tests and walk tests performed on the field on different types of small cells from multiple vendor equipment deployed on strand, rooftop and indoor locations. An average cell radius provided by strand mount small cell is approximately 300 meters. This is line of sight (LOS) radius and not a uniform radius of circular coverage as strand mount CBSDs usually deployed below clutter at 18 feet radiation point. The cell radius for line of sight coverage from an Attached mount (Rooftop) small cell is at least 1500 meters, diameter of cell coverage is 3000 meters. However, the overall non line of sight radius depends on height of attached mount CBSD and is relative to clutter and can vary from 300 meters to 700 meters radius (600 meters to 1400 meters diameter).

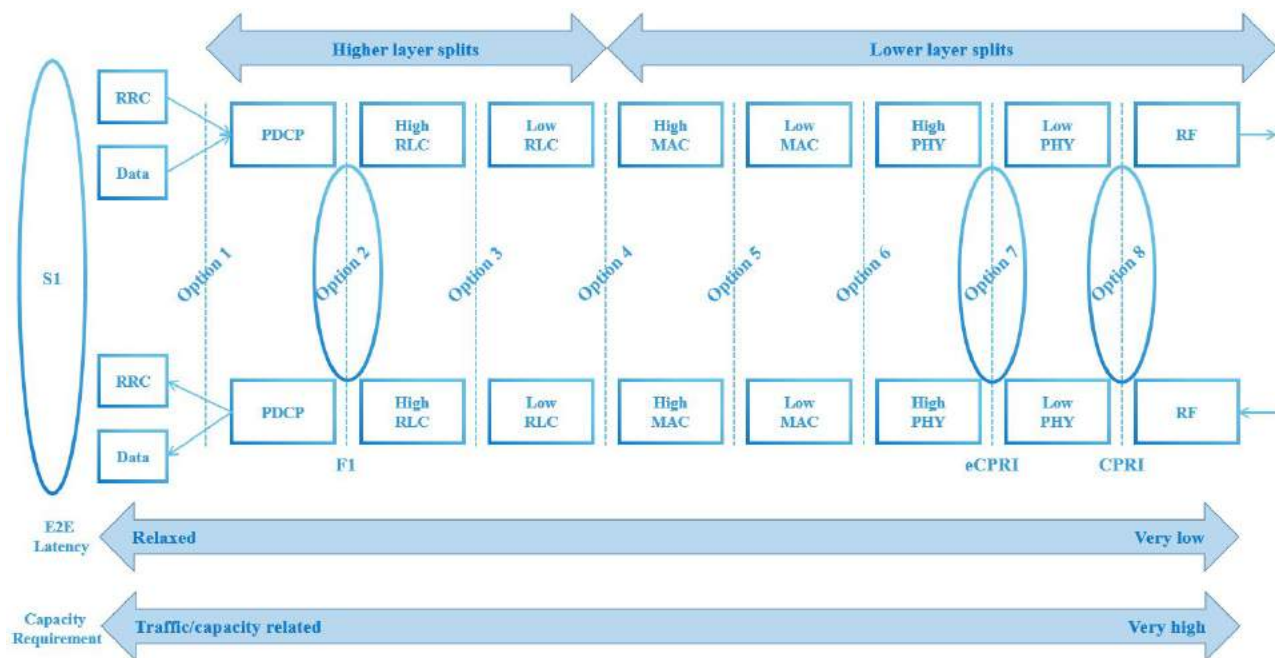
## 6. Introduction of 3GPP Virtual RAN Split Options

Virtual Radio Access Network (vRAN) is a new network deployment architecture that reduces the operational and capital expenditures of wireless network deployment. vRAN network architecture consists of mainly three parts, 1- Central Unit (CU), 2- Distributed Unit (DU) and 3- Radio Unit (RU). Each of these parts can execute different layer(s) of network communication protocol, and this is classified as ‘Split Options. In this part of white paper, we will explain what different split options are, and why Charter has selected split option 2 for DOCSIS based 5G CBRS wireless network deployment.

## 6.1. Split Options

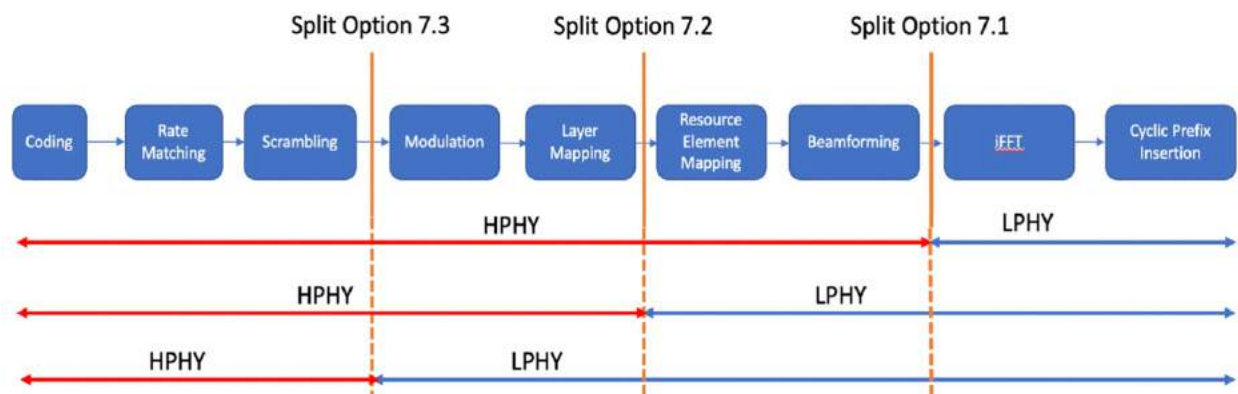
In vRAN deployment model, there are seven main protocol split models used to execute different parts of protocol layers at Distributed Unit (DU) and at Central Unit (CU). Here is the list of protocol split options used today.

- Option 1: RRC/PDCP split
- Option 2: PDCP/RLC split
- Option 3: Intra-RLC split based on a split of the RLC functionality such that the entire RLC is located in the central unit hosted in the cloud
- Option 5: Intra-MAC split
- Option 6: MAC/PHY split
- Option 7-1: This split where the IFFT and cyclic prefix insertion/removal performed at the remote radio unit (distributed unit) and IQ samples in frequency domain exchanged over the Fronthaul interface
- Option 7-2 & 7-2a: These splits have the additional benefit that pre-coding and digital beamforming, or parts thereof, are performed at the remote radio unit (distributed unit)
- Option 7-3: This split option considered only for the downlink, further reduces bandwidth requirements on the interface as coded user data exchanged before modulation.



**Figure 21 Protocol Split Options**

Each protocol layer has a number of processes that perform the tasks of that protocol layer. In addition, these processes represented by the processing blocks. In Physical layer (PHY) of LTE standard, there are nine processing blocks and split options 7.1, 7.2 and 7.3 created based on splits between processing blocks.

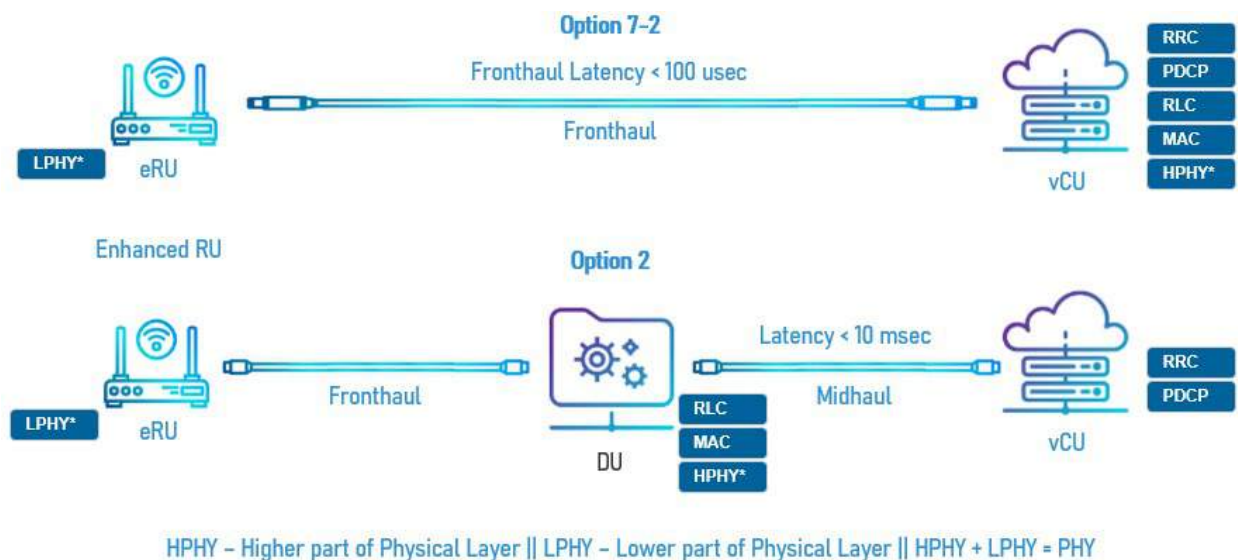


**Figure 22 Sub-splits for Split Option 7**

Processing blocks below split point are called Lower Physical Layer (LPHY) and processing blocks above split point is called Higher Physical Layer (HPHY). LPHY is executed in RU and HPHY is executed in DU.

## 6.2. Split Options and Latency Requirements


Latency requirements on the interface in general become tighter the further down in the protocol stack the split is, though one can classify two levels of latency requirements. For 3GPP, Split Option 1-3, where all RAN functionality that has to operate synchronously to the radio (i.e. in real time), the interface latency requirements are rather relaxed. Latency requirements are much more stringent for the other options where RAN split cuts in between synchronous functionalities. Recommendation to have Split Option 7-3 for CRAN deployment. Split Option 2 or Option 3 is preferred for deployment with DOCSIS since latency requirements is rather relaxed.



**Figure 23 Sub-splits Option 7-2 vs Split Option 2**

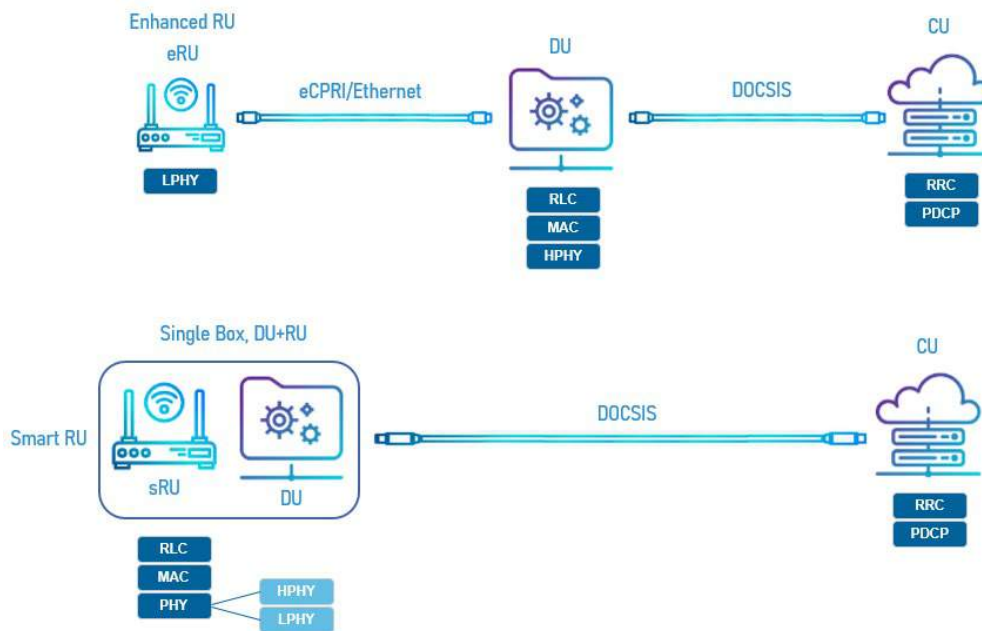
### 6.3. Split Option 2 vs Option 7-2

Split Option 2: As Charter, we selected split option 2 model for vRAN deployment over our widely available DOCSIS network. Split option two model has a separation at RLC layer, executes PHY, MAC and RLC layers at DU and RRC and PDCP layers at CU. Data transmission latency requirement for split option 2 is at most 10 msec. Any latency value of more than 10 msec. will degrade the quality of service delivered to our subscribers.



**Figure 24 Split Option 2**

Split option 2 has latency advantage over split option 7.2 in addition to the following advantages such as Mobility support. Handover between DUs will be similar to current intra-sector handover since DUs served with the same CU (same RRC layer). Link latency between DU and CU will affect the RRC messaging (required for HO) latency. No fiber requirements. Lower scheduling delay since scheduler runs in DU.




**Figure 25 Split Option 2 over DOCSIS**

Charter has selected combined RU and DU model for cable strand DOCSIS deployment. In this model, DU and RU is combined in a singled box executing PHY, MAC, HPHY (higher part of PHY layer) and LPHY (lower part of PHY layer) protocol layers.

Split Option 7-2: This model requires fiber deployment and we are considering using this split in markets where we have the required fiber coverage. Data latency requirement is 500 microseconds. Certain advantages of this split model are, but not limited to, Interference Management through Coordinated Multi-

Point (CoMP) and Inter-cell Interference management. SON features will benefit from central processing, and scheduling. Since eCPRI is not used, cost of CU will be lower.



**Figure 26 Split Option 7-2 over Fiber**

## 6.4. Conclusion

Charter has selected Split Option-2 for its vRAN based 5G CBRS wireless network deployment over DOCSIS because of its data transmission latency advantages. Also, since DOCSIS network provides advanced latency reduction features such as Bandwidth Report (BWR), Charter is ready to leverage these features to further ease of its 5G CBRS wireless network deployment on cable strand. We will keep monitoring the latest advancements in vRAN field, and latest releases of 3GPP standards together with oRAN initiative to fine tune our vRAN network deployment strategies.

## 7. DOCSIS Backhaul

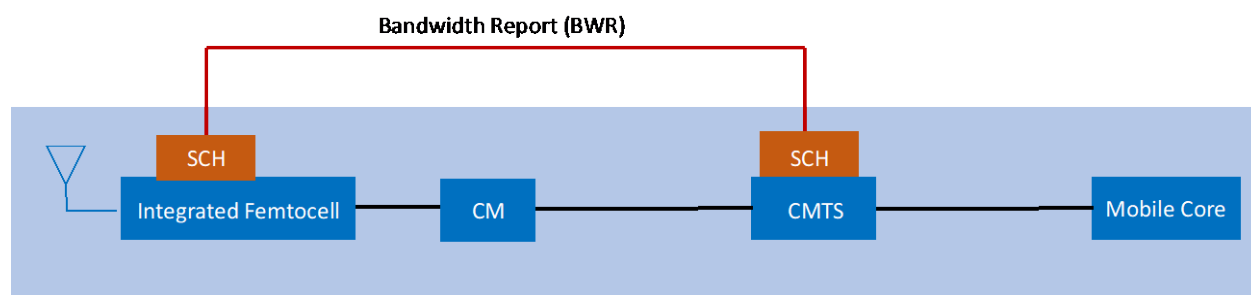
Near ubiquitous availability of Cable and DOCSIS assets in urban and suburban areas is one of the significant factors driving Charter's and more broadly cable industry's interest in Small Cell deployments for many years.

In anticipation of DOCSIS use as backhaul for wireless, the wireless technical leadership at Charter and CableLabs have been busy adding several critical and innovative features in DOCSIS.

In this section of the paper, we discuss these features.

### 7.1. Low Latency Xhaul (LLX)

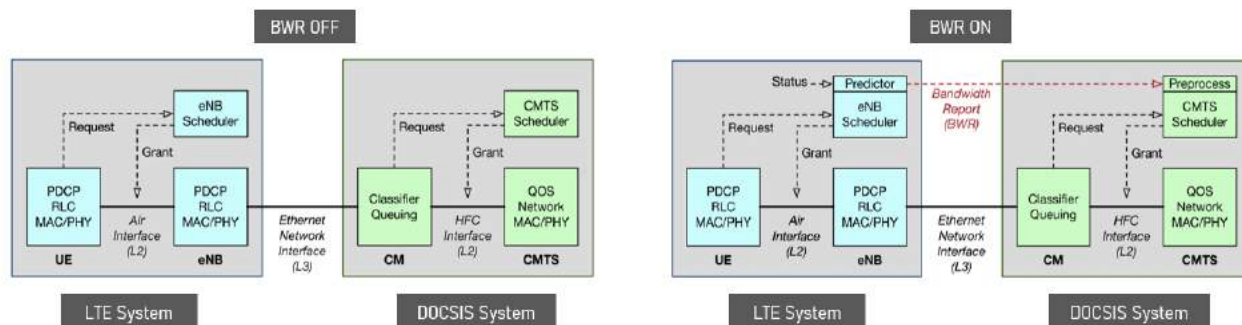
LLX aims to reduce the latency for mobile traffic on the DOCSIS network to as low as 1–2.5ms. It uses mobile and DOCSIS scheduler pipelining to achieve lower latency. CableLabs issued LLX specification in 2019 in a document titled Low Latency Mobile X-haul over DOCSIS® Technology. As shown in the picture, LLX introduces a new interface called Bandwidth Report (BWR) for exchange between the scheduling functions of the mobile and DOCSIS networks, which produces the effect of one pipelined system rather than two independent systems.



**Figure 27 Low Latency Xhaul**



The Bandwidth Report (BWR) is sent by the eNodeB scheduler to the DOCSIS scheduler and provides information about the amount of bytes eNodeB scheduler expects for some specific time in the future for the data transmission before the arrival of the actual traffic. Current LTE-DOCSIS systems do not have BWR and have a cumulative latency. The latencies of the eNodeB and CMTS systems are additive. Latency increases further with network congestion.



**Figure 28 Bandwidth Report OFF/ON**

For example, a 4G eNB or 5G gNB provides a future traffic profile through the BWR message, allowing the CMTS to make QoS and granting decisions earlier than it normally would. As a result, we can observe significant reduction in the variability in Jitter, improvement in average latency and smoothening of latency curve when BWR is ON. BWR works in 5G-NR similarly to the way it works in LTE for the majority of 5G traffic that employ larger slot sizes. 5G-NR URLLC (Ultra Reliable Low Latency Communication) latency requirements are 1-2ms. Using BWR the CMTS scheduler will have to predict the number of grants required per queue, verify this a few milliseconds later when the BWR arrives, and adjust as necessary.

## 7.2. Low Latency DOCSIS (LLD)

LLD is a CableLabs specification developed with vendors and operators. It focuses on lowering the latency by reducing queuing delay and media acquisition delay. The queuing delay is, lowered by using a different logical path for queue building vs. non-queue-building application traffic. In addition, the media acquisition delay is, lowered by using proactive scheduling mechanisms. LLD targets a reduction in round-trip latency in the DOCSIS network to below 5ms. It also targets a reduction in delay variation on the DOCSIS network by a factor of 100 compared to what is typical today. As a result, LLD promises to deliver significantly improved user experience for consumer applications, such as online gaming. Very importantly, the existing DOCSIS 3.1 equipment can include support for LLD with a software upgrade and without requiring hardware change.

## 7.3. Conclusion

Although Low Latency DOCSIS (LLD) and Low Latency X-Haul (LLX) have similar-sounding names, they are two different technologies with different objectives. While LLX reduces the latency for mobile user traffic on the DOCSIS network used as mobile X-haul (backhaul, mid-haul, or Fronthaul), the LLD reduces the latency in general on the DOCSIS network.

As the interest in Small Cell deployments is gaining momentum, operators are actively analyzing the LLX technology and inquiring the SmallCell vendors to support it. To operators' relief, the mobile eNB/gNB and the DOCSIS CMTS can add support for LLX via software upgrade without requiring hardware change.

As the Smallcell and CMTS vendor implementations of LLX become available in the next year, Charter looks forward to evaluating the technology in the lab and field.

## 8. Timing and Synchronization

Unlike Wi-Fi, the 4G LTE and 5G NR require stringent synchronization (Phase and Frequency) of wireless transmission to avoid interference between uplinks and downlinks. Since the CBRS is a shared band, the clock synchronization across basestations of the same and different operators is critical for full realization of the spectrum and avoid unwanted interference. As laid out in the table below, the synchronization requirements for TDD LTE and 5G NR are specially stringent.

**Table 13 Frequency and Phase Synchronization Requirements**

	<b>Frequency</b>	<b>Phase</b>
4G LTE TDD	$\pm 50$ ppb	$\pm 1.5$ $\mu$ s
5G NR TDD	$\pm 50$ ppb	$\pm 1.5$ $\mu$ s

These synchronization requirements are documented in 3GPP specifications – TS 36.133, TS 36.922, and TS 38.104.

As discussed earlier in this paper, Charter is pursuing both indoor and outdoor deployment of CBRS Smallcells. The acquisition of accurate phase and frequency for outdoor Smallcel deployment is rather straight forward and Charter is planning to use GPS.

On the other hand, the acquisition of accurate phase and frequency for indoor wireless deployment is much more complex and requires evaluation of multiple options.

As shown in the table, there are a number of options for timing and synchronization. For outdoor deployments, GPS is the most widely used timing source. However, for indoor applications such as Femtocell, the combination of PTP and DTP ranks higher on the list and is carefully studied and tested.

Table 14 Synchronization Types - Advantages and Disadvantages

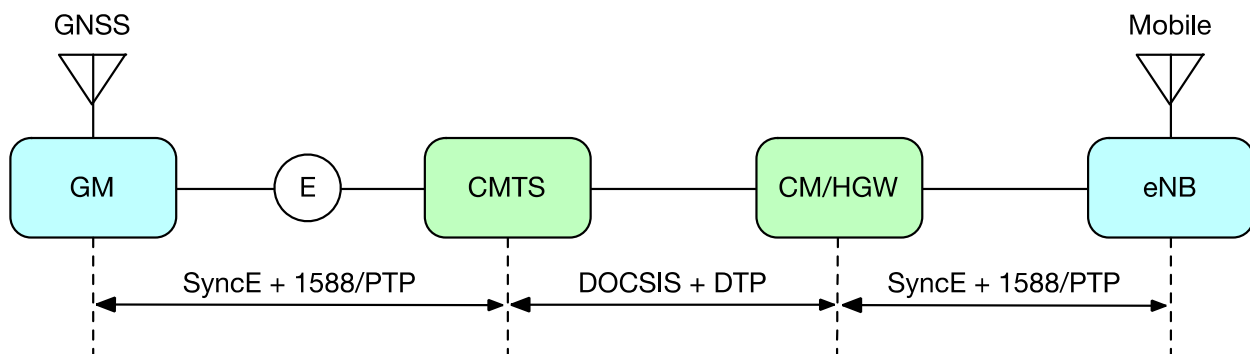
	Advantages	Disadvantages
DOCSIS Timing Protocol (DTP) with PTP	<ul style="list-style-type: none"> <li>• Supports LTE TDD and 5G timing precision requirements [2]</li> <li>• Timing from our Charter owned and operated network</li> <li>• CableLabs standard promoted by cable vendors</li> </ul>	<ul style="list-style-type: none"> <li>• Requires significant changes to DOCSIS infrastructure, including hardware upgrade to CM</li> <li>• Grand Master clocks in each headend</li> <li>• Regular network calibrations may be required</li> </ul>
Over-The-Top PTP	<ul style="list-style-type: none"> <li>• No upgrades to DOCSIS network required</li> </ul>	<ul style="list-style-type: none"> <li>• Timing synchronization not precise enough for TDD LTE even with DOCSIS QoS. (5-10 millisecond range)</li> <li>• Performance is negatively impacted with network loading and uplink packet delay variation (uplink BW limited)</li> </ul>
Network Listen/Macro Sniffing	<ul style="list-style-type: none"> <li>• No upgrades to DOCSIS network required</li> </ul>	<ul style="list-style-type: none"> <li>• Reliance on Macro network for timing</li> <li>• Availability everywhere is an issue</li> <li>• Out-of-band listen requires dedicated radio – additional cost &amp; more space</li> </ul>
Global Positioning System (GPS)	<ul style="list-style-type: none"> <li>• No upgrades to DOCSIS network required</li> <li>• Supports LTE TDD timing precision requirements</li> </ul>	<ul style="list-style-type: none"> <li>• Receive challenges indoors, susceptible to jamming</li> <li>• Placement not in the control of the operator</li> <li>• Installation and operation cost external antennas</li> </ul>
Over the top Network Time Protocol (NTP)	<ul style="list-style-type: none"> <li>• No upgrades to DOCSIS network required</li> </ul>	<ul style="list-style-type: none"> <li>• Timing synchronization not precise enough (100 millisecond) even with dedicated QoS on DOCSIS</li> </ul>
TV Broadcast Listen	<ul style="list-style-type: none"> <li>• No upgrades to DOCSIS network required</li> </ul>	<ul style="list-style-type: none"> <li>• Need a receiver for TV broadcast</li> <li>• Femtocell must know location – own &amp; TV tower</li> </ul>

### 8.1. DOCSIS Timing Protocol (DTP)

DTP is part of CableLabs' DOCSIS 3.1 specifications and further enhanced in CableLabs CM-SP-SYNC specification. DTP carries the IEEE 1588v2 protocol over the DOCSIS network without the influence of jitter from network buffering. Additionally, DTP accounts for plant and path asymmetry that over-the-top

PTP may not. Furthermore, DTP provides timestamp and determines the downstream timing offset, which PTP uses to calculate precise timing and synchronization.

The picture below shows the application of DTP and PTP when DOCSIS network is used as backhaul for the mobile traffic. As shown, the DTP used between the CMTS and CM and PTP is between the grandmaster clock and CMTS and between the CM and eNB. CMTS and CM essentially convert PTP to DTP on the DOCSIS link. Sync E makes the frequency synchronization performance of PTP better. In case of loss of the master clock, Sync E also helps eNB keep clock running longer in holdover mode without losing precision.



**Figure 29 DOCSIS Timing Protocol**

DTP is a new protocol and is currently not widely implemented or deployed. However, there is a keen interest in it from cable operators exploring indoor Small Cell deployments. Operators need a reliable source of precise timing and synchronization for both 4G TDD LTE and 5G NR deployments.

## 8.2. Assisted GPS

For indoor deployments, Charter is also evaluating “Assisted” GPS technology, which claims to provide accurate phase and frequency in challenging RF environments (e.g., building basement). In an “Assisted” GPS enabled system, DOCSIS backhaul could be used to feed the small amount of data carried by the satellite signals. This process makes it unnecessary for the “Assisted” GPS receiver to demodulate the data, allowing the system to provide accurate phase and the frequency at significantly low GPS signal levels.

## 8.3. Conclusion

The engineering teams at Charter are actively testing both PTP/DTP and “assisted” GPS technologies in collaboration with CableLabs and Cisco. We hope to publish results from our study in the next Cable-Tec-Expo for the benefit of the industry.

## 9. Wrap-Up

The rollout of a wireless network is a significant challenge. The wireless engineering team at Charter has dedicated the past three years to understanding the characteristics which are important to a successful wireless deployment. In this paper, we have covered the important deployment scenarios which should be understood for every deployment – Strand Mount (Outdoor Aerial Strand), Attached Mount (Outdoor, non DOCSIS), and Indoor Deployments. Strand Mount utilizes all of the advantages of the HFC/DOCSIS network such as power availability, wide coverage, and quick deployments. This does need to be countered with the additional power load on the network. The Attached Mount deployments are utilized for strategic applications such as hot spots which can not be reached via the HFC network. The most important consideration of an Attached Mount is power and permitting which can eat up months of time getting a unit deployed. Key to success though is ultimately the indoor coverage. A drawback of CBRS is the EIRP power limits imposed by the FCC. This significantly reduces the ability of a 3.5 GHz signal to enter most building structures; therefore, it is important to have a strong indoor coverage plan. Charter has found a significant advantage with an Indoor/Outdoor strategy where two sectors cover the street outside of a building and the third sector is for indoor coverage.

It is important to understand the RF properties of CBRS. The 3.5 GHz signal is in a sweet spot for new wireless spectrum. It has the advantages of small form factor versus lower frequency spectrum while still getting reasonable coverage versus mmWave spectrum. From our research, RSRP is critical to monitor and manage. An RSRP value which drops below -110 / -115 dB causes a significant drop in throughput. It is important to rely on accurate 3D modeling to make sure a reasonable RSRP is maintained throughout the network. With the timeline of a wireless network deployment, it is important to consider 5G in your decision. Most significant deployments of a wireless network will start in 2-3 years which will give time for the initial higher cost of 5G to align with present day 4G deployment costs. Organizations like ORAN are helping drive 5G costs down quicker through standards and interoperability.

With the deployment of a DOCSIS network, it is important to consider Split 2 versus Split 7.2. Our research and testing has shown that Split 2 or a full gNB is a more practical option for a DOCSIS network primarily due to the amount of capacity required to support Split 7.2. Latency is a challenge, but it can be managed for a Split 7.2 configuration via adoption of technologies such as Low Latency Xhaul (LLX) and Low Latency DOCSIS (LLD).

For indoor and dense urban environments, timing will be a challenge. There is research and new technology in development focused on supporting better timing in both of these environments. Our current focus is on DOCSIS Timing Protocol and Assisted GPS. Both technologies have proven to support a better timing signal indoors.

It is hard to put three years of research into a single paper. The focus of this paper has been on the technologies which we feel have the strongest impact on a new wireless network on a modern day HFC plant. Best of luck with your wireless deployment.

# Abbreviations

3GPP	3 <sup>rd</sup> Generation Partnership Project
AOI	Area of Interest
AWS	Advanced Wireless Services
BH	Busyhour
bps	bits per second
BWR	Bandwidth Report
CBRS	Citizens Broadband Radio Service
CBSD	Citizens Broadband Radio Service Device
CE	Cell Edge
CM	cable modem
CMTS	cable modem termination system
CN	Cell Near
CPAS	Coordinated Periodic Activates among SASs
CRAN	Centralized Radio Access Network
CU	Central Unit
CW	continuous wave
DL	downlink
DLS	Decision Logic System
DOCSIS	Data Over Cable Service Interface Specifications
DPA	Dynamic Protection Area
DSDS	Dual SIM Dual Standby
DTP	DOCSIS Timing Protocol
DU	Distributed Unit
E-DPA	ESC Dynamic Protection Area
EIRP	Effective Isotropic Radiated Power
ESC	Environmental Sensing Capability
EXZ	Exclusion Zone
FCC	Federal Communications Commission
FEC	forward error correction
FSS	Fixed Satellite Station
FWA	Fixed Wireless Access
GAA	General Authorized Access
GHz	gigahertz
GM	Grand master
GPS	Global Positioning System
GSM	Global System for Mobile Communications
GWBL	Grandfathered Wireless Broadband Licensee
HD	high definition
HFC	hybrid fiber-coax
HO	Handover
HPHY	Higher Physical layer
Hz	hertz
IFFT	Inverse Fast Fourier Transform
LLD	Low Latency DOCSIS

LLX	Low Latency Xhaul
LOS	line of sight
LPHY	Lower Physical layer
LTE	Long Term Evolution
MAC	media access control
Mbps	Megabits per second
MHz	megahertz
MIMO	Multiple Input Multiple Output
MVNO	Mobile Virtual Network Operator
NTIA	National Telecommunications and Information Administration
NTP	Network Time Protocol
OOB	Out of Band
PAL	Priority Access License
PCAST	President's Council of Advisors on Science and Technology
PDCP	Packet Data Convergence Protocol
P-DPA	Portal Dynamic Protection Area
PDSCH	Physical Downlink Shared Channel
PHY	Physical layer
PPS	pulse per second
PSD	power spectral density
PUSCH	Physical Uplink Shared Channel
QAM	Quadrature Amplitude Modulation
QOS	Quality of Service
QPSK	Quadrature Phase Shift Keying
RF	radio frequency
RLC	Radio Link Control
RRC	Radio Resource Control
RSRP	Reference Signals Received Power
RU	Radio Unit
SAS	Spectrum Access System
SCTE	Society of Cable Telecommunications Engineers
SINR	Signal to Noise and Interference Ratio
SMB	small medium business
TT&C	telemetry, tracking and command
UE	user equipment
UL	uplink
vRAN	Virtual Radio Access Network

## Bibliography & References

*Code of Federal Regulations, Title 47, Part 96*  
*The Wireless Innovation Forum*  
*CableLabs® Specifications*

# Wi-Fi 6 And Wi-Fi 6E Are Building The Foundation For New Home Applications

A Technical Paper prepared for SCTE•ISBE by

**Massinissa Lalam, Ph.D**

Sr. Expert on Wireless Technology  
Sagemcom  
250 Route de l'Empereur; 92500 Rueil-Malmaison, France  
+33 1 57 61 13 41  
massinissa.lalam@sagemcom.com

**Kamal Koshy**

Sr. Director Wireless Engineering  
Charter Communications  
Denver, CO  
kamal.koshy@charter.com

**Xavier Briard**

SVP Broadband Solutions  
Sagemcom  
Dallas, TX  
xavier.briard@sagemcom.com



# Table of Contents

Title	Page Number
1. Introduction.....	4
2. Evolution of Wi-Fi Technology .....	4
2.1. Early Wi-Fi Generations .....	5
2.2. Generational Wi-Fi .....	5
2.3. Peak Throughput Evolution.....	6
3. Wi-Fi 6 Technology Overview .....	7
3.1. Geared For Multi-User Transmissions .....	7
3.2. Toward an AP-centric Channel Access .....	8
3.2.1. Classical Channel Access Mechanim .....	8
3.2.2. 11ax Trigger Frame Channel Access .....	10
4. Wi-Fi Challenges at Home .....	10
4.1. Entertainment .....	11
4.2. Smart Home .....	11
4.3. Internet access .....	11
5. How Wi-Fi 6E Will Change Experience in Dense Residential Areas .....	11
5.1. Experience in Legacy Band .....	11
5.1.1. Inteference free setup .....	12
5.1.2. OBSS Setup.....	13
5.2. The 6 GHz Experience .....	15
6. Conclusion.....	16
Abbreviations .....	17
Bibliography & References.....	17

## List of Figures

Title	Page Number
Figure 1 - Throughput increase over the Wi-Fi generation .....	6
Figure 2 - DL/UL OFDMA.....	7
Figure 3 - DL/UL OFDMA efficiency for small packets delivery.....	7
Figure 4 - DL/UL MU-MIMO .....	8
Figure 5 - Channel access example (pre-11ax).....	9
Figure 6 - Example of a Trigger Frame exchange for UL-OFDMA transmission.....	10
Figure 7 - Interference Free Setup.....	12
Figure 8 - Wi-Fi 5 Throughput Evolution (Interference Free Setup) .....	13
Figure 9 - Wi-Fi 6 Throughput Evolution (Interference Free Setup) .....	13
Figure 10 - Wi-Fi 5 Throughput Evolution (OBSS Setup) .....	14
Figure 11 - Wi-Fi 6 Throughput Evolution (OBSS Setup) .....	14
Figure 12 - 6 GHz Spectrum (US).....	15
Figure 13 - Wi-Fi 6 Throughput Evolution (6 GHz-like Setup) .....	16

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Wi-Fi 5 vs Wi-Fi 6 Average Throughput (Interference Free Setup).....	12
Table 2 - Wi-Fi 5 vs Wi-Fi 6 Average Throughput (OBSS Setup).....	14
Table 3 - Wi-Fi 6 Average Throughput (6 GHz-like Setup).....	15

## 1. Introduction

Years after years, Wi-Fi has become the de-facto way to access Internet at home. Based on the IEEE 802.11 standard family, the Wi-Fi technology has experienced several major PHY/MAC updates during the last 20 years, each one defining a generation on its own. Not so long ago, the general public was still not aware of those gaps and just called “Wi-Fi” any one of them. The generational naming introduced by the Wi-Fi Alliance in 2018 was seen a better way to advertise each new release instead of relying on 802.11 amendments name.

Built upon the IEEE 802.11 ax amendment, Wi-Fi technology is currently at its sixth coined “Wi-Fi 6” which operates on the unlicensed 2.4 and 5 GHz ISM band with an extension into the 6 GHz territory called Wi-Fi 6E.

Each Wi-Fi generation brought an improvement in peak throughput: from the original 11Mbps back in 1999, Wi-Fi 6/6E now offers up to 9.6 Gbps. However, Wi-Fi 6/6E main target was not peak throughput increase but improved efficiency in dense environments.

Indeed, the success of Wi-Fi, due to its low cost, ease of use and performance led to an explosion of Wi-Fi devices, exhibiting the limits of previous generations in terms of congestion and channel access in such scenarios.

Prior to Wi-Fi 6, access points (APs) and stations (STAs) were contending to access the medium with similar priority. More and more end-devices being deployed, resulting congestion has led to a degraded experience in crowded places like multi-dwelling units (MDUs).

In this paper, we will briefly present the evolution of IEEE 802.11/Wi-Fi technology and its quest for more throughput. We will then present the main features of Wi-Fi 6 and the change of paradigm it brought to the table with a more AP-centric channel access. We will then discuss how Wi-Fi has changed the residential environment and the main uses cases it needs to address. We will then present how Wi-Fi 6 can change user experience at home today and how the opening of the 6 GHz band will drastically change the user experience tomorrow. With three times more spectrum available, results achieved in clean environment could be truly representative of the end user experience even in dense environment.

## 2. Evolution of Wi-Fi Technology

Wi-Fi technology, which brand is owned by the Wi-Fi Alliance (WFA), is based on IEEE 802.11 standard. While IEEE 802.11 develops the technology, the WFA aims to promote it and creates as such interoperability testing and certification programs to ensure the end-users of a certain level of performance and compatibility. While independent, both organizations are tightly linked in the way Wi-Fi became so successful nowadays. General public knows the brand “Wi-Fi” but is generally not aware of the main generations it has experienced in its life, partly because Wi-Fi has always been backward compatible with its previous generations. An equipment bought in 1999 could work today with the same performance even in the most advanced Wi-Fi network. The generational naming introduced in 2018 by the WFA with numbers is a step to better differentiate the Wi-Fi technologies. Before addressing Wi-Fi 6/6E, we will recall what are the different Wi-Fi generations and their quest to always increase the throughput.

## 2.1. Early Wi-Fi Generations

The first commercial success of the IEEE 802.11 technology was relying on the 11b amendment back in 1999. Operating on the 2.4 GHz industrial, scientific and medical (ISM) band, this technology was based on direct sequence spread spectrum (DSSS) and complementary code keying (CCK) modulations and offered up to 11 megabits per second (Mbps) at the physical (PHY) layer on channel occupying 22 MHz. Its inclusion in laptop's CPU helped democratizing what can be considered as the first generation of Wi-Fi.

While development started at the same time as 11b, products based on 11a amendment came later. This technology was solely operating on the 5 GHz ISM band where more spectrum was available. It was relying on orthogonal frequency division multiplexing (OFDM), which divides the bandwidth in smaller subcarriers for a transmission in the frequency domain. Back in that day 11a could delivered a throughput up to 54 Mbps at PHY layer through the use of higher order modulations (64 quadrature amplitude modulation (QAM)) and a channel bandwidth of 20 MHz. This can be considered as the second generation of Wi-Fi.

The will to port most of 11a innovations introduced in the 5 GHz ISM band into the 2.4 GHz ISM band gave birth to the 11g amendment in 2003. The same OFDM and modulations as 11a were introduced, while keeping backward compatibility with 11b devices. Only the 20 MHz channel bandwidth from 11a was kept. As such 11g offered up to 54 Mbps at the PHY layer with 20 MHz channel and could be considered as the third generation of Wi-Fi.


## 2.2. Generational Wi-Fi

In order to increase even more the throughput and as computation capability was more affordable, IEEE 802.11 worked on a new amendment which would be defined for both 2.4 and 5 GHz ISM band. With 11n, peak PHY throughput was increased by doubling the bandwidth (from 20 MHz to 40 MHz, an option viable mainly in the 5 GHz ISM band which is larger) and introducing multiple-input multiple-output (MIMO) concept which allow the use of up to four antennas. If channel conditions permit, spatial multiplexing which is a mode of MIMO transmission allow the transmission on the same channel of different parallel streams, increasing the throughput. 11n could then offer a throughput up to 600 Mbps at the PHY layer. 11n was published in 2009 by IEEE 802.11, but some years before people could buy products supporting a draft version of 11n pushed by the WFA to respond to the public demand. 11n is the fourth generation of Wi-Fi and named as Wi-Fi 4 by the WFA.

As the most efficient way to increase throughput is to use more spectrum, the next IEEE 802.11 amendment was only operating at the 5 GHz ISM band where more than two 20-MHz channels could be aggregated. 11ac amendment use all possible dimensions available to increase its peak throughput: higher bandwidth (up to 160 MHz), higher modulation order (up to 256 QAM), higher MIMO (up to 8 antennas). A full feature 11ac could offer up to 6.9 gigabits per second (Gbps) at the PHY layer, while the vast majority of gateways were shipped with up to 4 antennas. 11ac was published in 2013 by IEEE 802.11. The WFA used only a subset of 11ac features for its wave 1 certification based on a stable 11ac draft, later completed with additional ones for its wave 2 program. The use of larger channel was possible because spectrum was made available in the 5 GHz band. However, all 160 MHz channels available in this band overlap with incumbents (e.g. radar) which have primary access. The constraint made on the devices to detect those signature and move away in case of positive detection (process often called dynamic frequency selection (DFS)) is such that in practice only 80 MHz channels are really used in residential deployments. 11ac is the fifth generation of Wi-Fi and named as Wi-Fi 5 by the WFA.

After the completion of 11ac, IEEE 802.11 group was trying to address a problem which was more and more visible as Wi-Fi technology was more and more successful. Originally thought as a cable replacement, the issue of density of devices and congestion associated to it was not really part of the design. With the 11ax amendment, IEEE was targeting to improve efficiency in dense deployment scenarios, with a throughput increase measured this time not at PHY layer but above the medium control access (MAC) layer, i.e., at the application layer, directly related to the user experience in such deployment. Since no work have been carried out on the 2.4 GHz ISM band since 2009, it was decided from the beginning to define 11ax for both the 2.4 and 5 GHz band, with a later extension to the 6 GHz band. 11ax brings a lot of changes which we will detail in a later section, but it also increases the throughput by using four times more subcarriers within the same spectrum (guard band can be reduced while guard interval (typical of OFDM modulation) is proportionally smaller) and up to 1024 QAM modulation. Therefore, a full 11ax solution can deliver up to 9.6 Gbps. 11ax amendment is expected to be published in 2021 by IEEE 802.11, but products are already available. WFA release 1 certification is ready since mid-2019 based on a subset of 11ax features from a stable draft. This certification only covers the 2.4 and 5 GHz band and it is expected that the 6 GHz band will be covered soon. 11ax is the sixth generation of Wi-Fi called Wi-Fi 6, while Wi-Fi 6E is used to indicate support of the 6 GHz band.

### 2.3. Peak Throughput Evolution



**Figure 1 - Throughput increase over the Wi-Fi generation**

As shown from Figure 1, Wi-Fi technology, like any successful wireless technology, has always known peak throughput increase through its various generations. Wi-Fi 6/6E is particular in the sense that this metric was not its primary objective. Wi-Fi 6/6E was developed to improve efficiency in dense environment like MDUs or stadium, where the limit come from too much devices trying to access a limited spectrum available at the same time.

### 3. Wi-Fi 6 Technology Overview

Wi-Fi 6 introduced many features to address the high density problem. In the following, we will detail the most important one when facing a residential unmanaged deployment.

#### 3.1. Geared For Multi-User Transmissions

When too many stations are doing traffic at the same time, one solution to reduce congestion is to group them to either send them data (downlink) or receive data from them (uplink).

Wi-Fi 5 AP could already send data up to 4 STAs at the same time with a method relying on spatial separation called DL MU-MIMO. This technique is useful to exploit the asymmetry in antenna configurations between an AP (equipped generally with 4 antennas) and the STA (usually no more than 2 antennas) to send more streams to different STAs. However, this approach is heavily dependent on the radio conditions, topology and signal processing at the STA side.

To cope with this, Wi-Fi 6 introduced OFDMA as a way to group several STAs together in the frequency domain. Supported in both DL and UL directions, this method allocates group of (contiguous) subcarriers to different STAs.




Figure 2 - DL/UL OFDMA

This technique is particularly efficient when packets to transmit are small, reducing the air time needed to send them.





Figure 3 - DL/UL OFDMA efficiency for small packets delivery

If traffic is heavy with larger packet size, OFDMA could be less efficient due to the fact that the bandwidth is divided among the users. But Wi-Fi 6 can still use DL MU-MIMO transmissions to deal with such use case. Later Wi-Fi 6 products are also expected to be able to use this technique in the reverse direction as well with UL MU-MIMO.



**Figure 4 - DL/UL MU-MIMO**

In Wi-Fi 6, scheduling at the AP side takes an even more important role as the AP is now in charge of DL and UL transmission. To make the STAs transmit in the uplink when the AP wants to, a new trigger mechanism is introduced which allow a more AP-centric medium access as we will discuss below.

## **3.2. Toward an AP-centric Channel Access**

Before explaining how the trigger principle works and how it can improve efficiency at the network level, it is worth recalling how channel access used to work prior to Wi-Fi 6.

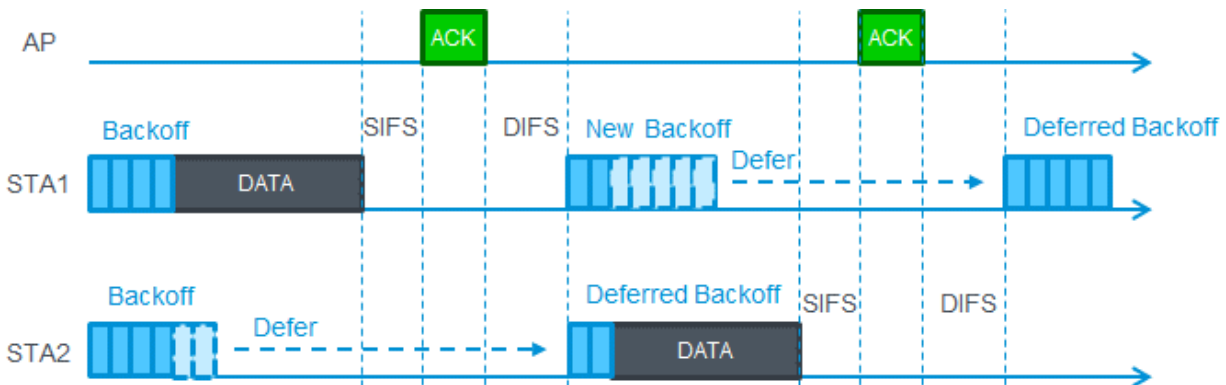
### **3.2.1. Classical Channel Access Mechanim**

During the history of Wi-Fi products, the channel access mechanism has been mainly updated from the distributed coordination function (DCF) access originally defined in the first IEEE802.11 standard from 1997<sup>1</sup>. Based on carrier sense multiple access with collision avoidance (CSMA/CA), each device wanting to transmit shall listen first to the medium to detect any energy above a given threshold. If none is detected, then it can transmit. Such principle is also known as the listen before talk (LBT) principle, terminology often used in regulation. Of course some refinements were added to this simple principle to avoid systematic collision between two devices such as the use of a backoff counter and an exponential contention window.

A device, being it a STA or an AP, shall pick a random value between 0 and the actual contention window size and store it in a counter. After sensing the medium for a fixed duration (DIFS), if the medium is idle (i.e. no energy detected above the threshold), then the device can decrement the counter by one for each time slot (lasting 16μs or 9μs) it senses the medium idle. When the counter reaches zero, the device can transmit. If energy is detected while the counter was not zero, then the device should stop until no more energy is detected and resume the decount after the fixed duration sensing (DIFS). If the frame sent is not acknowledged, then the device shall pick a new value but this time it shall double the

<sup>1</sup> Another mechanism based on polling, called point coordinated function (PCF), was also introduced but never encountered success through commercial products.

contention window size to reduce the probability of collision. Figure 5 shows an example of two STAs contending for the medium.



**Figure 5 - Channel access example (pre-11ax)**

Additionally to this energy detection mechanism, Wi-Fi relies also on a preamble detection PHY mechanism. Preamble is a part of each Wi-Fi frame which indicates the duration of the given frame. By successfully decoding it, the sensing device knows when it can resume its sensing if it wanted to transmit. This process is further refined by the fact that if the sensing device decodes the complete frame, then the MAC header contains information on the duration on the ongoing exchange (though the network allocation vector (NAV) field), i.e.; including duration of the next frames to be sent which helps the sensing device even more.

Modifications have been added to this channel access mechanism, in particular through the 802.11e amendment from IEEE 802.11 and the wireless multimedia (WMM) certification from WFA, like aggregation of MAC and service packets (AMPDU, AMSDU respectively) to increase the throughput of the support of quality of service (QoS). QoS is achieved by having different parameters for the contention window size minimum and maximum values and the fixed duration used for sensing (AIFS) leading to an enhanced distributed channel access (EDCA).

More details can be found in the relevant standards and specifications, but the main principle pre-11ax was that all devices should operate using this principle. This means that within a Wi-Fi network, or basic service set (BSS), the AP and all its associated STAs contend equally for the medium. Note that when QoS is engaged, the AP may have a slight advantage in terms of access with default parameters for:

- best effort category: with a maximum contention windows size lower than the one to be used by the STAs,
- voice/video categories: with a lower fixed duration used for sensing.

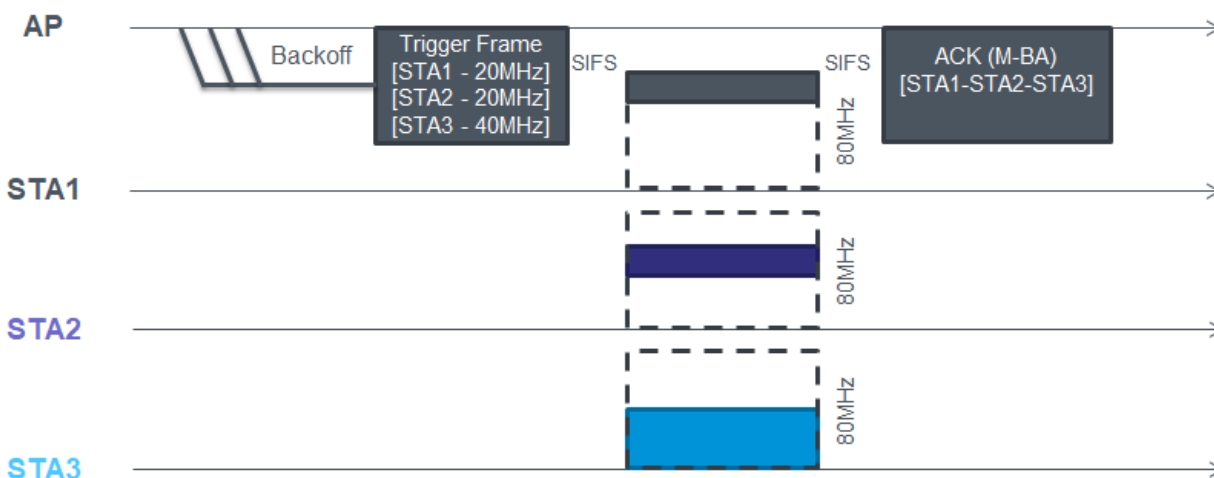
With more and more devices being deployed at home, AP is competing with more and more STAs increasing collision over the air and congestion. This becomes a greater problem when other Wi-Fi networks are deployed in the vicinity on the same channel. The networks, called overlapping BSSs (OBSS) have APs but also STAs connected to them. One can easily understand that such crowded environment could lead to bad user experience: all devices fighting for the medium with (almost) the same priority.



### 3.2.2. 11ax Trigger Frame Channel Access

To mitigate this effect, Wi-Fi 6 decided to shift the paradigm of the channel access to make it more AP-centric instead of device-centric for the uplink direction. Instead of having all STAs contending for the medium, only their AP will do so. For example, if we have 2 BSSs with  $N$  STAs associated to each one, instead of having  $2N+2$  devices trying to access the medium, only the 2 APs will contend for both downlink (DL) and uplink (UL) transmissions.

To coordinate multi-user transmissions in uplink the trigger frame (TF) was introduced. This frame, sent by the AP, schedules the next UL transmissions of all addressed STAs. Figure 6 shows an UL-OFDMA transmission on a 80 MHz BSS where STA1 is allocated the first 20 MHz, STA2 is allocated the second 20 MHz and STA3 is allocated the remaining 40 MHz. The AP can acknowledge the reception of all uplink transmission with a single ACK frame (called M-BA). In this case instead of three stations contending for the medium, only the AP will send its trigger frame



**Figure 6 - Example of a Trigger Frame exchange for UL-OFDMA transmission**

Once engaged at least once in such uplink procedure, the STAs will refrain to access the channel autonomously leaving only their AP in charge of gaining access to the medium for them.

Of course this mechanism can only be applied by Wi-Fi 6 STAs. But as more of them are entering the market right now, legacy devices will also see some benefits: the newer STAs will not contend for the medium but only their APs will. In fact, by choosing to address dense deployment, Wi-Fi 6 true potential will only be achieved if enough Wi-Fi 6 capable device are in the market. So while the trend is in favor of it, we have to remember that the vast majority of the devices at home are still Wi-Fi 4/5-only capable.

## 4. Wi-Fi Challenges at Home

Wi-Fi as a technology has become synonymous to internet both in residential and enterprise environments. From its inception in the late 1990s, Wi-Fi performance has improved by orders of magnitude as shown previously. The use cases supported by Wi-Fi has also significantly changed over the years. In a typical residential environment, Wi-Fi may support an average of 15 devices (2020), and up to 50 devices in a tech centric residence. The use cases supported by now includes internet access, gaming devices, IoT devices, mobile devices with varied used cases. The following sections detail some of the key use cases that Wi-Fi as a technology needs to address

## **4.1. Entertainment**

With streaming services that are available, home entertainment has moved to wireless from the earlier fixed media devices like DVD, CD etc. Supporting robust streaming services require dedicated bandwidth to media consumption devices including TVs, Roku, media players and mobile devices used as media consumption devices. Simultaneous media consumption with multiple devices also needs to be supported by Wi-Fi technology. Additional support for newer use devices that support augmented reality(AR) and virtual reality (VR) is also a requirement for Wi-Fi technology.

## **4.2. Smart Home**

Smart Home include variety of devices including video devices (video doorbell, cameras) and IoT sensors (Wi-Fi and other technologies). Smart video devices require sufficient bandwidth to support video upload/download. IoT sensors based on Wi-Fi require technology support for long battery life and reliable communication at the edges of the residence.

## **4.3. Internet access**

The requirements of internet access to homes has gone up significantly over the years. As mentioned earlier, the number of devices in many US homes is 15 and in many cases approach 50 devices. The increase in devices places significant burden on the network to provide great connectivity to devices. In addition, due to the shared nature of the access medium in Wi-Fi, additional technologies have to be developed to support the increase in number of devices. Multi-dwelling units present additional challenges since the number of devices that are in proximity can be significantly higher than in standalone residences. New technologies and best practices have to be developed to support internet access in multi-dwelling units.

The recent introduction of Wi-Fi 6 should help addressing the previous use cases while the opening of the 6 GHz band should alleviate the most critical issue in MDUs which is congestion due to lack of spectrum.

# **5. How Wi-Fi 6E Will Change Experience in Dense Residential Areas**


## **5.1. Experience in Legacy Band**

Where Wi-Fi 6 really shines, to the point of improving the user experience by a factor 4, is really in the heavy congested environments with lots of users sending small packets. The trigger frames sent by the AP to schedule uplink transmissions of its associated STAs dramatically reduces the congestion otherwise observed with the classical EDCA mechanism where all STAs contend for the medium.

However, you do not need a dense deployment to see the benefit of Wi-Fi 6 which may be not throughput related. To demonstrate it, we set ourselves in a clean environment (i.e. no neighboring networks). We set up an 11ax 4x4 Wi-Fi 6AP configured on a 80 MHz channel, similar to what is available in the 5 GHz band without having to rely on an DFS mechanism. We associated to it :

- two 2x2 160 MHz-capable Wi-Fi 6 laptops, and
- one 4x4 Wi-Fi 6 160 MHz-capable router set as a station (think of it as an extender for instance).

The devices are close to the AP with equivalent received power. Figure 7 shows our setup which is an interference free one.



**Figure 7 - Interference Free Setup**

To evaluate any benefit of Wi-Fi 6 in such scenario, we run throughput tests (using IxChariot) when setting the AP first in Wi-Fi 5 mode then in Wi-Fi 6 mode. For both runs, we use the following traffic profiles at the same time:

- a UDP uplink transmission with a throughput target of 250 Mbps for each laptop,
- a UDP downlink transmission with a throughput target of 500 Mbps for the router.

UDP is used to show to see Wi-Fi 6 benefits at MAC/PHY layer without any recovering protocol from upper layers. We use the same quality of service for all traffics (Best Effort).

#### **5.1.1. Inteference free setup**

Table 1 shows the average throughput obtained after 1 minute of traffic.


**Table 1 - Wi-Fi 5 vs Wi-Fi 6 Average Throughput (Interference Free Setup)**

Device	Traffic Direction	Throughput Target (Mbps)	Wi-Fi 5 (Mbps)	Wi-Fi 6 (Mbps)
Laptop #1	Uplink	250	238	209
Laptop #2	Uplink	250	156	217
Router	Downlink	500	491	494


We see that in both mode downlink traffic is almost the same, the 500 Mbps target is almost honored for Wi-Fi 5 and Wi-Fi 6. However, the uplink results show a different story. In Wi-Fi 5 mod, one laptop is getting significantly more throughput (238 Mbps) than the other (156 Mbps). The radio conditions being the same, this demonstrates that some STA implementations could be more aggressive than others in their channel access mechanism (with use of frame bursting for instance). In Wi-Fi 6 however, both STAs have

equivalent uplink throughput (209 Mbps vs 217 Mbps) which is controlled by the AP. Both do not reach the target throughput (250 Mbps) though.

The full story is shown by the throughput curves during the time which are given in Figure 8 and Figure 9 for Wi-Fi 5 and Wi-Fi 6, respectively.



**Figure 8 - Wi-Fi 5 Throughput Evolution (Interference Free Setup)**



**Figure 9 - Wi-Fi 6 Throughput Evolution (Interference Free Setup)**

All Wi-Fi 5 curves are more impacted by the lack of coordination between the three sources of traffic which all contend to access the medium. Wi-Fi 6 curves are less prone of such variation since only the AP controls the traffic by triggering the two laptops and instructing them to use UL OFDMA transmissions while it serves also the third device in downlink. One can see that both STA curves (blue and red) are almost the same leading to a better fairness. Ultimately, this will also result in better latency for all three devices with a service delivery more stable thanks to Wi-Fi 6.

### **5.1.2. OBSS Setup**


In an MDU type of deployment however, you are usually not alone on an 80 MHz channel. To see the effect, we introduced an interfering network on the same channel as our setup (OBSS) but at a reasonable distance from it. In the lab, we reduced the transmit power of such interferer to simulate distance. On this

setup we run a continuous TCP full-buffer traffic (iperf) to create a Wi-Fi noise to which our setup has to defer (mainly due to preamble detection).


We ran the same experiment in both Wi-Fi 5 and Wi-Fi 6 to check the effect of OBSS in way which is closer to what people are experiencing. Results are given Table 2, Figure 10 and Figure 11.

**Table 2 - Wi-Fi 5 vs Wi-Fi 6 Average Throughput (OBSS Setup)**

Device	Traffic Direction	Throughput Target (Mbps)	Wi-Fi 5 (Mbps)	Wi-Fi 6 (Mbps)
Laptop #1	Uplink	250	132	168
Laptop #2	Uplink	250	68	176
Router	Downlink	500	479	486



**Figure 10 - Wi-Fi 5 Throughput Evolution (OBSS Setup)**




**Figure 11 - Wi-Fi 6 Throughput Evolution (OBSS Setup)**

As one can see, both setups are affected by the OBSS and none could reach the throughput target, but the Wi-Fi 6 one offers more stability to both uplink transmissions with fair access of the spectrum to both of them.

## 5.2. The 6 GHz Experience

With the recent FCC decision to open the 6 GHz band to wireless devices, more than 1.2 GHz band will be made available for Wi-Fi 6E in the US. It is envisaged that residential AP deployment may choose to use low-power indoor model to avoid the access to a geolocation database for automatic frequency selection (AFC). Under such conditions, the whole 6 GHz band is available without any constraints à-la DFS. To put it simply, seven 160 MHz channels will be available for a Wi-Fi 6E AP to choose from, greatly reducing the risk of OBSS really affecting the performance.




**Figure 12 - 6 GHz Spectrum (US)**

In dense residential deployment, most of the congestion/collision will then come from inside the BSS itself. Higher frequency means more wall attenuation and with 7 channels to choose, overlapping should not be experienced in the near future on this band. With no legacy devices present, STAs will most likely choose to be controlled by the AP for a better user experience.

It is envisaged to have 160 MHz channel bandwidth as the norm in 6 GHz band. If we take our previous interference free setup to make it operate in a 6 GHz-like environment, then we can setup our AP with a 160 MHz channel bandwidth and repeat our measurements. Table 3 shows that in this 6 GHz-like environment targets are reached on average for the three devices (less than 2% deviation) while Figure 13 shows the throughput evolution and its stability

**Table 3 - Wi-Fi 6 Average Throughput (6 GHz-like Setup)**

Device	Traffic Direction	Throughput Target (Mbps)	Wi-Fi 6E (Mbps)
Laptop #1	Uplink	250	168
Laptop #2	Uplink	250	176
Router	Downlink	500	486



**Figure 13 - Wi-Fi 6 Throughput Evolution (6 GHz-like Setup)**

6 GHz opening is also considered in Europe, but on the lower 500 MHz band first. Recent decision from Ofcom UK to open it by end of this year provides three 160 MHz channel to Wi-Fi 6E devices for operation. While this is less than what the US are opening right now, the key advantage of Wi-Fi 6 are still there to improve experience at home.

## 6. Conclusion

In this paper we reviewed the evolution of Wi-Fi technologies over the years and the challenge it should address in dense residential scenarios. If the introduction of Wi-Fi 6 technology certainly helps in providing unique features to improve user experience in dense deployment, the wireless systems are still bounded by the spectrum available to them. Fortunately, the recent opening of the 6 GHz band in the US, paving the way to others countries (Europe, South Korea, ...), allows Wi-Fi 6E to really be able to change user experience at home, even in dense areas since Wi-Fi 6 key features are more effective in Wi-Fi 6 only ecosystem. Indeed, IEEE 802.11 operation on 6Ghz band is restricted to Wi-Fi 6E devices only (i.e. no Wi-Fi 4/5 or even 11b legacy to take into considerations).

## Abbreviations

AFC	automatic frequency selection
AP	access point
A-MPDU	aggregated MPDU
A-MSDU	aggregated MSDU
BSS	basic service set
CCK	complementary code keying
CSMA/CA	carrier sense multiple access with collision avoidance
DCF	distributed coordinated function
DFS	dynamic frequency selection
DL	downlink
DSSS	Direct sequence spread spectrum
EDCA	enhanced distributed channel access
Gbps	gigabits per second
GHz	giga hertz
IEEE	Institute of Electrical and Electronics Engineers
IFS	inter-frame spacing
ISM	industrial, scientific and medical
LBT	listen before talk
MAC	medium access control
Mbps	megabits per second
MDU	multi-dwelling unit
MIMO	multiple input multiple output
MPDU	MAC protocol data unit
MSDU	MAC service data unit
MU	multi-user
OBSS	overlapping basic service set
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiplexing access
PCF	point distributed function
QAM	quadrature amplitude modulation
QoS	quality of service
STA	station
TCP	transmission control protocol
TF	trigger frame
UDP	user datagram protocol
UL	uplink
WFA	Wi-Fi Alliance

## Bibliography & References

*IEEE P802.11ax - IEEE Draft Standard for Information Technology -- Telecommunications and Information Exchange Between Systems Local and Metropolitan Area Networks -- Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*



*Amendment Enhancements for High Efficiency WLAN*, Institute of Electrical and Electronics Engineers, Draft 6.0, 2020

*Wi-Fi Alliance® introduces Wi-Fi 6*, Wi-Fi Alliance, press release, <https://www.wi-fi.org/news-events/newsroom/wi-fi-alliance-introduces-wi-fi-6> [last access 28/07/2020]

*802.11-2016 - IEEE Standard for Information technology--Telecommunications and information exchange between systems Local and metropolitan area networks--Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, Institute of Electrical and Electronics Engineers, 2016

*FCC Opens 6 GHz Band to Wi-Fi and Other Unlicensed Uses*, Federal Communications Commission, <https://www.fcc.gov/document/fcc-opens-6-ghz-band-wi-fi-and-other-unlicensed-uses-0> [last access 28/07/2020]

*Statement: Improving spectrum access for Wi-Fi – spectrum use in the 5 and 6 GHz bands*, Ofcom UK, press release, <https://www.ofcom.org.uk/consultations-and-statements/category-2/improving-spectrum-access-for-wi-fi> [last access 28/07/2020]

# Using Big Data To Fine-Tune The Nation's Largest Public Wi-Fi Network

A Technical Paper prepared for SCTE•ISBE by

**Eli Baruch**

VP, Solution Engineering  
CommScope  
2 Logan Sq. Suite 1810 Philadelphia, PA  
4049932001  
Eli.baruch@commscope.com

Indira Paudel

Manager XfinityWiFi Operation  
Comcast  
1800 Arch St. Philadelphia, PA  
2155102268  
Indira\_Paudel@comcast.com

# Table of Contents

Title	Page Number
1. Introduction .....	4
2. Testing and Methodology .....	4
2.1. Manual vs. Automated Methodology .....	6
2.1. Current Sources of Data for Analytic system .....	7
3. Defining and Measuring User Experience .....	8
3.1. Use Case for XfinityWiFi network .....	9
3.2. Key Performance Indicator .....	9
4. Identifying Bad WIFI Client Experience (Bad WCX) .....	13
4.1. Criteria's of Bad WIFI Client Experience (Bad WCX) .....	13
4.2. Types of Bad WCX events .....	13
4.2.1. Example of identifying type of Bad WCX .....	14
4.1. Correlating Bad WIFI Client Experience with WIFI SNMP Data .....	14
5. Current Results and Findings .....	17
5.1. Baseline state .....	17
5.1. Current state .....	18
6. Conclusion .....	21
Abbreviations .....	22

## List of Figures

Title	Page Number
Figure 1 Breaking up the problem – optimization zones (Primary Zones are highlighted) .....	5
Figure 2 High-Level view of the usage of XfinityWiFi network .....	6
Figure 3 High-Level description of a big data analytics platform .....	8
Figure 4 DHCP DORA 4-way exchange .....	10
Figure 5 TCP 3-way handshake .....	10
Figure 6 Average Channel Health Statistics .....	15
Figure 7 Average Calculated Channel Throughput .....	15
Figure 8 Outliers of Channel Health Statistics .....	16
Figure 9 5GHz Radio Data MPDU distribution across MCS .....	16
Figure 10 Percentage of sessions with Bad WCX out of all client sessions – Baseline state .....	17
Figure 11 Bad WCX breakdown per type – Baseline state .....	17
Figure 12 Bad WCX distribution per Radio – Baseline state .....	18
Figure 13 Bad WCX distribution per AP Model – Baseline state .....	18
Figure 14 Bad WCX distribution per AP Model – Baseline vs. Current .....	19
Figure 15 Percentage of sessions with Bad WCX out of all client sessions and improvement .....	19
Figure 16 Bad WCX distribution per AP Model – Current state .....	20
Figure 17 Percentage of sessions with Bad WCX out of all client sessions and improvement – OG1600s .....	20

## List of Tables

Title	Page Number
Table 1 Sources of Data .....	8

Table 2 Example SNMP information elements provided by the Access Point.....	11
Table 3 Example of classification to the type of Bad WCX.....	14

# 1. Introduction

How does one go about optimizing the most extensive public Wi-Fi network in the nation? That was the task we were asked to perform. Comcast XfinityWiFi network is a vast Wi-Fi overlay network, with close to 21M access points. These Access Points are deployed outdoors, inside small and medium businesses, and in subscriber's homes. The purpose of this ongoing project is to improve the customer experience of the XfinityWiFi network and increase the overall traffic usage, and increase the data offload of Xfinity Mobile users over the XfinityWiFi network.

The fundamental impact when customer faces a bad experience on Wi-Fi is that they may decide to turn off Wi-Fi on their mobile devices and may use valuable LTE data. While they may do so while outdoors, our data shows that sometimes these customers will continue to use LTE data while in their own homes and would not connect to their home Wireless Gateway provided by Comcast.

Some of the challenges facing the optimization of XfinityWiFi network, apart from the sheer size, were the fact that no central or distributed controller manages and coordinates the system from a Wi-Fi perspective. The Access Points are not aware of each other and have no communication path between them. The network was never designed for complete coverage or even to allow seamless roaming between Access Points. Also, over 98% of the network comprises of residential Access Points that act as a Home Hotspot (HHS). We had minimal control over what Wi-Fi configuration changes can be applied to these devices.

The first step we took was to break the optimization of a nationwide network into smaller, more manageable zones. We took several of such zones and used them to prove the methods, processes, and tools. We broke the process into multiple iterations, using traditional RF planning, measurements, and predictive tools to test after each iteration. Thus, establishing a baseline and iterative improvements as we progressed.

In parallel, the CommScope team developed a Big Data pipeline to ingest data from multiple sources of measurement; Multiple Network telemetry data, Access Point telemetry, enriched with geographic information data. This Big Data Analytics platform is turning discrete sources of data into powerful insights and recommendations.

This paper will describe in detail the methods, processes, and tools we developed to address this challenge and the results observed so far.

## 2. Testing and Methodology

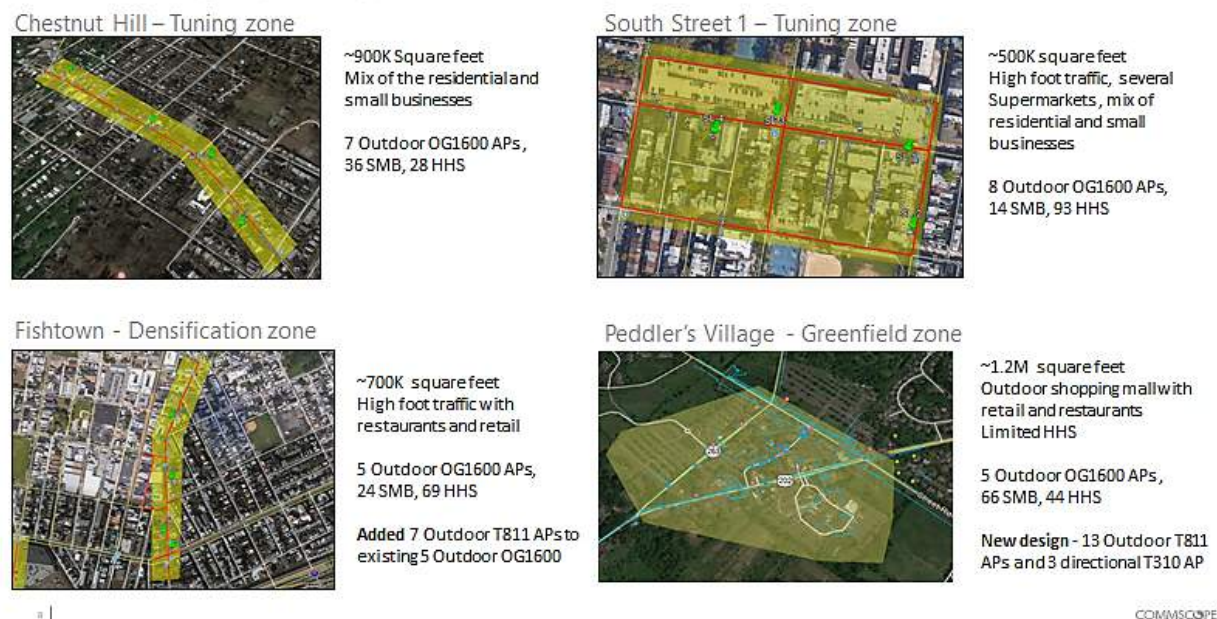
Working with Comcast's XfinityWiFi team, we agreed to address the problem of optimization of this nationwide network by breaking it into smaller geographical areas. The logic behind this approach is the fact that it is local radio interferences and environmental conditions that impact Wi-Fi Access Points (AP).

Comcast team selected several outdoor areas, with different characteristics, all had a mix of outdoor, and indoor APs; the indoor APs were a combination of the residential home hotspot (HHS) and business hotspot (SMB). Each of the selected areas had a Primary and Secondary zone. Manual testing were conducted in the primary zones.

We used different approaches for the optimization per designated primary zone:

1. Tuning Zone – No physical change to an existing outdoor network. (e.g., No adding of new Access point). Only modifications allowed were configuration changes.
2. Densification Zone – Allow changes to the existing outdoor network by adding new outdoor Access Points. Include configuration changes.
3. Greenfield Zone – A complete redesign of the outdoor access point deployment. Include configuration changes.

We established a Test Methodology and perform identical RF testing in each of the Primary zones. Third-party testers were using a combination of mobile clients (iOS and Android) to perform walking and stationary tests. Conducting standard RF measurement and Over the Air packet capture tools (Ekahau, Omnipeek, Accuver (Android only)). They measured RSSI, SNR, Noise Floor, and Throughput test while walking and in stationary locations. The tests establish baseline conditions in each Primary zone. The same testing repeated after each iteration and produced a detailed test report. A team of CommScope experts analyzed the raw data, and results and suggested modifications for the next iteration.



**Figure 1 Breaking up the problem – optimization zones (Primary Zones are highlighted)**


The main tools used in the iteration process, based on the limitations of the network design, were fundamental Wi-Fi configuration changes such as RF channel planning, channel width, transmit power, and modification of Wi-Fi cell size. In the Densification and Greenfield zones, we were able to add AP to ensure better coverage, and in the Greenfield even created an overlap between the outdoor AP.

Note, no changes were made to any of the indoor hotspots, residential or SMB.

Below is a visual description of the activity (connected clients and tonnage of usage) of the network during several months before the COVID-19 pandemic.

Sum of Active Users by Date and type

type (Blank) Home Outdoor SMB



Top 5 AP Models by total active users

AP Model	Sum of Active Users
OG1600A	18586677
TG1682G	3779329
DPC3941T	2904284
CGM4140COM	2850665
DPC3941B	2144322

Sum of Active Users by VLAN



daily Active Users by type

type (Blank) Home Outdoor SMB



dailyTonnage by type

type (Blank) Home Outdoor SMB



Figure 2 High-Level view of the usage of XfinityWiFi network

## 2.1. Manual vs. Automated Methodology

The problem with the method and testing methodology described above is that it does not scale. It relies on testers walking the areas, and a manual process that cannot be scaled nationwide.

Our proposed solution was to augment manual RF testing with measurements collected from the network. This includes using Netscout network probes reports and collecting relevant data from outdoor AP. This approach obtained data from network assets, independent of testers walking the areas. And thus, has the advantage of looking at real customer data vs. test client's data. To collect and make use of all this data, we created a Big Data Analytics platform and pipeline to ingest, and analyze different sources of data, cross-correlate multiple measurements into a comprehensive view. And use the manual testing process to “verify” our automated network results and insight.

The main goal of the big data approach and tools is to show how we will be able to answer the following questions:

- How to optimize and automate the process of “making XfinityWifi network better”?
- How to identify where and how the currently deployed network should be augmented?
- Help identify where Comcast should deploy the new XfinityWifi network/APs to offer more offload opportunities to Xfinity Mobile customers.
- Identify Bad Client Experience events, characterize and group them into subcategories
- Show some correlation between Netscout records and Wi-Fi RF (SNMP logs) records relating to the Bad Client Experience events and subcategories.

Several Machine Learning methods and algorithms were used on the existing data. Clustering, anomaly detection and co-occurrence were used to analyze the existing network; Identify areas of over coverage, under coverage, patterns of client movement across the network deployed in the zone and more.

## **2.1. Current Sources of Data for Analytic system**

All the traffic in the XfinityWiFi network traverses the network over softGRE tunnels that are established between the Access Point (AP) and the Wireless Access Gateway (WAG). By placing Netscout probes in the network close to the WAG, which monitors all TCP traffic inside the tunnel; The Comcast team was able to collect and report in 5-minute intervals several TCP measurements.

In addition to the TCP traffic reports generated by the network probes, we collected measurements from the outdoor APs that represent the Wireless performance. The outdoor Wireless Access Points reported a vast number of Wi-Fi RF and network statistics in 15-minute intervals.

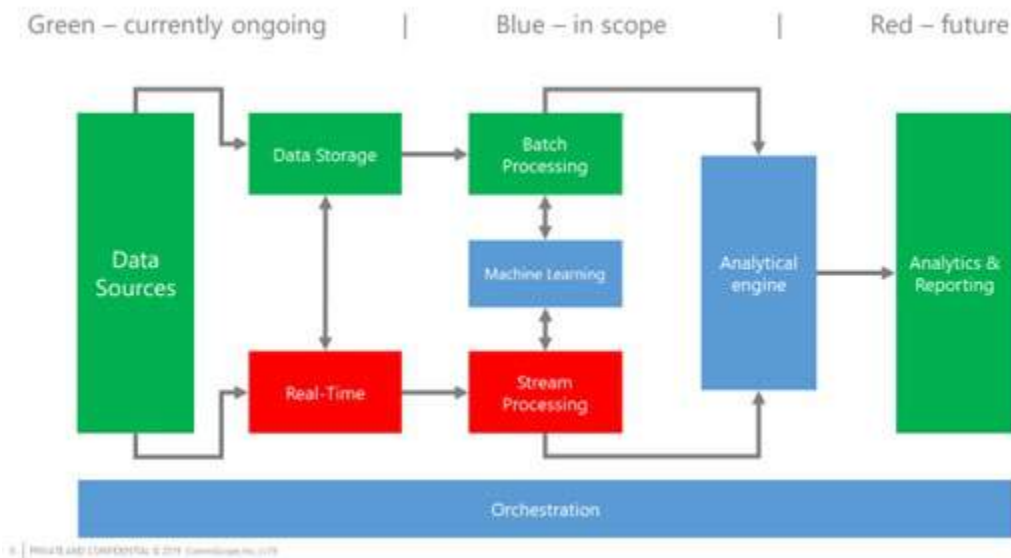
We enriched the information provided by Netscout and the APs with geographic information about the AP location, either by GPS coordinates or other GIS data.



**Table 1 Sources of Data**

Devices	Information	Source
Device location	Geo Location	Google Earth map (KMZ) and GPS coordinates
Outside plant location	Geo Location	Google Earth map (KMZ)
Ruckus Outdoor AP	Usage and Statistics	Netscout and SCI/Druid
OG1600 Outdoor AP	Usage and Statistics	Netscout and SNMP logs
Indoor HHS	Usage and Statistics	Netscout only
Indoor SMB	Usage and Statistics	Netscout only

The above different sources of data accumulated to 250+ days of data, 4000+ geographic locations, 10k-100k records per site per day. 10 TB of compressed data/month. 4 disparate data sources (csv, json, mib, log, kmz). With different types of triggers, event-based, 5 minutes, 15 minutes, and daily time windows.



**Figure 3 High-Level description of a big data analytics platform**

### 3. Defining and Measuring User Experience

One of the main objectives of the project is to improve the overall Wi-Fi customer experience. But how do we define good customer experience? Customers usually do not know or care about any of the RF parameters that govern wireless communication. Often, the user's subjective assessment of the quality of the service boils down to a few essential characteristics.

- Connection reliability, is the connection easy? Is it fast? Is it repeatable?
- Latency, do I get a snappy response from the network?
- Throughput, is there an adequate minimum throughput guaranteed?

Our challenge is how to quantify these subjective assessments, by translating them into parameters we can measure and control.

From the network side, we can measure many aspects that enforces the above assessment of quality:

- Packet loss and packet retransmission
- Airtime utilization
- Throughput
- Latency
- Consistency

Assuming we can measure, how do we know what constitutes Good or rather Bad Customer Experience? For that, we need to understand the use case and expectations from the service this network provides.

### **3.1. Use Case for XfinityWiFi network**

There are fundamentally two distinct use cases for the XfinityWiFi network. It may serve as a well-known “Guest Network” for visitors that come into your home or business. In this case, residential or business owners do not have to create a unique guest network and remember the credentials that would allow the visitor to join the network. All the access points broadcast the same SSID name (aka. Network name), and all Xfinity customers can join them for free. These users tend to be stationary users with mobile or handheld devices and laptops.


The second use case serves customers that are mainly outdoor; these are pedestrians with mobile or handheld devices. As such, the XfinityWiFi Network aims to create a continuous coverage in designated outdoor areas. And serve these pedestrian users’ data applications such as web browsing, video and audio streaming, IP voice application such as WhatsApp, Facetime, etc.

By having a good understanding of the above use cases, we can provide a good foundation for defining the Key Performance Indicator (KPI) that will be measured, their expected values and acceptable range.

### **3.2. Key Performance Indicator**


Netscout data provided by the Comcast team collected and reported in 5-minute intervals the following TCP measurements:

DHCP Time- Measure how quickly a customer can get an IP from the network by recording the time it takes to complete a DHCP complete DORA (discovery, offer, request, and acknowledgment) exchange.



**Figure 4 DHCP DORA 4-way exchange**

Client Latency- Measures the initial TCP 3-way handshake between client and server and calculates the client-side latency (time elapsed between the Syn-Ack and Ack)



**Figure 5 TCP 3-way handshake**

Average Throughput- records the TCP throughput in each direction (Upstream and Downstream) within a 5-minute measuring interval. Throughput is calculated by dividing the sum of bytes in each direction by the number of milliseconds of activity in the corresponding direction.

TCP Retransmission- Number and percentage of TCP packets that were retransmitted in each direction i.e, from the client towards the server, and from the server towards the client.

While the first two measurements (DHCP Time and Client Latency) only represent a brief snapshot in time, the other two are ongoing and provide much more detailed measurement of all active TCP sessions between the client and the network.

The Average Throughput report is also used to sum the total of traffic (Tonnage) that traverse the network over a period.

In addition to the TCP traffic reports generated by the network probes, we collected measurements from the outdoor APs that represent the Wireless performance.

The outdoor Wireless Access Points reported a vast number of RF and network statistics in 15-minute intervals. Below are some of the relevant measurements collected.

**Table 2 Example SNMP information elements provided by the Access Point**

Group	Information Element	Description
Client information	Client IP address	
	Client MAC Address	for all clients associated with the AP
	Client State	Authenticated, Associated, Authorized, Secure or Unsecure
	Client RSSI	the last Received Signal Strength Indicator (RSSI) of the wireless client device.
	Min RSSI	the minimum RSSI of the wireless client device through the 15-minute interval.
	Max RSSI	the maximum RSSI of the wireless client device through the 15-minute interval.
	Client Last Transmit Rate	the transmit rate of the last MPDU sent to the client
	Client Last Receive Rate	the receive rate of the last MPDU sent from the client
Active Wi-Fi Channel quality statistics	Min Noise Floor	the minimum noise floor (dBm) on the serving channel across the measurement interval.
	Max Noise Floor	the maximum noise floor(dBm) on the serving channel across the measurement interval.
	Median Noise Floor	the median noise floor (dBm) on the serving channel across the measurement interval
	Activity Factor	percentage of airtime the radio was actively utilizing the channel across the measurement interval
	Channel Utilization	percentage of time the medium was utilized on the channel across the measurement interval

	Retransmission	percentage of packets that had to be retransmitted during the measurement interval
	Busy by Other Radio	percentage of each second that the transmitter is busy by another Radio
	Average Tx Mod Rate	a weighted average of the Tx modulation rate for data MPDU in Mbps
	Average Rx Mod Rate	a weighted average of the Rx modulation rate for data MPDU in Mbps
Radio based statistics	Radio Tx OK	number of successfully transmitted MSDU (data MPDU)
	Radio Tx Fail	the number of hardware Tx errors and excessive retries for MSDU (data MPDU).
	Radio Rx OK	number of Ethernet frames forwarded to stack by wifi MAC
	Radio Rx Fail	number of data MPDU missed by sequence number gap
	Tx MCS	the number of Data MPDU transmitted for each MCS (modulation and coding scheme).
	Rx MCS	the number of Data MPDU received for each MCS (modulation and coding scheme).
SSID based statistics	SSID Tx OK Bytes	per SSID number of bytes in successfully transmitted MSDU (data MPDU).
	SSID Tx OK	per SSID number of successfully transmitted MSDU (data MPDU).
	SSID Rx OK Bytes	per SSID number of bytes in Ethernet frames forwarded to stack by wifi MAC.
	SSID Rx OK	per SSID number of Ethernet frames forwarded
	SSID Successful Auth	per SSID number of wifi clients successfully authenticated
	SSID Auth Failure	per SSID number of wifi client's authentication Failures

## 4. Identifying Bad WIFI Client Experience (Bad WCX)

To provide a better user experience, we had to define what does “good user experience” means. And therefore, what are the bad cases we wish to identify, eliminate, or correct.

DHCP Time and Client Latency only represent a brief snapshot in time; however, they provide insight into the Connection reliability and the perceived latency. We were aiming for a fast Authentication (802.1x or Open SSID), and fast DHCP response (less than 1Sec). These parameters of Authentication and DHCP response, are entirely under the network operator control. As a rule of thumb, we also defined expectable latency as less than 150ms.

In order to measure latency, throughput, and reliability on an ongoing basis, we used the following criteria’s as “good user experience”:

1. Less than 4% of TCP packet retransmission,  
Measurements conducted by Comcast identified 4% as the threshold where clients start noticing network delay and slowness.
2. A minimum downstream throughput in the coverage areas,  
We were aiming at HD Video quality inside cell coverage (>5Mb/s) and SD Video quality at cell edge (>3Mb/s).  
Note, Outdoor AP may serve up to 100 active clients per Radio, 50 per SSID.

Evidence of Bad WIFI Client Experience (Bad WCX) can be gleaned from records of a high percentage of TCP packets retransmission. Another proof of Bad WCX is slow average throughput for lengthy sessions. We should look for a combination of indications from the Netscout reports of the same client MAC at the same 5 min bin.

### 4.1. Criteria’s of Bad WIFI Client Experience (Bad WCX)

1. Percentage of Server Retries  $\geq 4\%$  AND Number of packets retries  $> 2$  (packets)
2. Percentage of Client Retries  $\geq 4\%$  AND Number of packets retries  $> 2$  (packets)
3. Video sessions that are longer than 2 minutes AND average downstream throughput less than 3Mbps

### 4.2. Types of Bad WCX events

Looking at the total Bad WCX events based on the above criteria, we can group them into different types of events, which may require different resolutions.

1. Moving Client
  - records identify client on Multiple AP during the 5 min bin
2. Client on a Single AP, Single VAP/VLAN.
3. Client on a Single AP, Multiple VAPs/VLAN.

Note – in order to create this classification of events; we would need to add within the 5min bin all records of the client that experienced a Bad WCX within that 5min bin.

### 4.2.1. Example of identifying type of Bad WCX

As seen in Table 3 below. At 11:45 The client 48:BF:6B:92:FB:D8 is connected to an OG1600 with AP\_MAC\_Address of a8:9f:ec:da:b2:6a, experienced a Bad WCX while connecting to the 5GHz xfinitywifi VAP. However, the same client also registered on the 2.4GHz VAP during the same time. And exchanged most of the traffic on this VAP. This is a case of **Client on a Single AP, Multiple VAPs/VLAN**

At 13:50 and 14:05 the same client experienced 2 Bad WCX events. These are cases of the **client on a Single AP, Single VAP/VLAN**.

At 16:10 we see the same client registered on multiple AP, experiencing 2 Bad WCX events. On TG34862G with AP\_MAC\_Address 88:ef:16:dd:b0:01. And on TG862G AP\_MAC\_Address cc:a4:62:af:68:d4. This is a case of **Moving Client**

**Table 3 Example of classification to the type of Bad WCX**

Time Line	5 Min	AP			From TCP_retri es_From Server Packets	%TCP Server Retransm ission	TCP_retri es_From Client Packets	%TCP To Server Client Retransm ission
		MAC Address	Model	VLAN				
11:45	6/15/2020 11:45	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_5GHz_OG	612	3	23	13.04%
11:45	6/15/2020 11:45	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	6	1,218	0.49%
12:30	6/15/2020 12:30	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	12	290	4.14%
13:20	6/15/2020 13:20	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_5GHz_OG	612	12	3,773	0.32%
13:20	6/15/2020 13:20	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	8	88	9.09%
13:50	6/15/2020 13:50	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	22	187	11.76%
14:05	6/15/2020 14:05	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	3	27	11.11%
14:50	6/15/2020 14:50	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	26	473	5.50%
15:00	6/15/2020 15:00	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	42	871	4.82%
15:05	6/15/2020 15:05	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	40	335	11.94%
15:10	6/15/2020 15:10	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	31	377	8.22%
15:30	6/15/2020 15:30	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_5GHz_OG	612	5	3,211	0.16%
15:35	6/15/2020 15:35	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	8	130	6.15%
15:55	6/15/2020 15:55	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	58	850	6.82%
16:00	6/15/2020 16:00	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	21	506	4.15%
16:05	6/15/2020 16:05	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_5GHz_OG	612	9	980	0.92%
16:05	6/15/2020 16:05	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	22	157	14.01%
16:10	6/15/2020 16:10	88:ef:16:dd:b0:01	TG3482G	xfinitywifi_2_4GHz_HHS_BWG	102	6	81	7.41%
16:10	6/15/2020 16:10	cc:a4:62:af:68:d4	TG862G	xfinitywifi_2_4GHz_HHS_BWG	102	3	9	33.33%
17:10	6/15/2020 17:10	80:b2:34:3d:52:e5	DPC3941T	xfinitywifi_5GHz_HHS_BWG	103	3	19	15.79%
17:10	6/15/2020 17:10	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	6	29	20.69%
17:15	6/15/2020 17:15	a8:9f:ec:da:b2:6a	OG1600A	xfinitywifi_2_4GHz_OG	602	9	86	10.47%

### 4.1. Correlating Bad WIFI Client Experience with WIFI SNMP Data

Utilizing telemetry, currently in the form of SNMP polls or Ruckus SmartCell Insight (SCI) reports. Allows us to make measurements without the need to have testers on the ground. Furthermore, collecting data in small intervals, over a long period enable us to identify trends in behaviors. And avoid any transient events that may impact test results of a specific day.

The main trends we were looking for were Channel Utilization, Activity Factor, Wi-Fi Retransmissions, and calculated Average Modulation Rate.




Figure 6 Average Channel Health Statistics





Figure 7 Average Calculated Channel Throughput

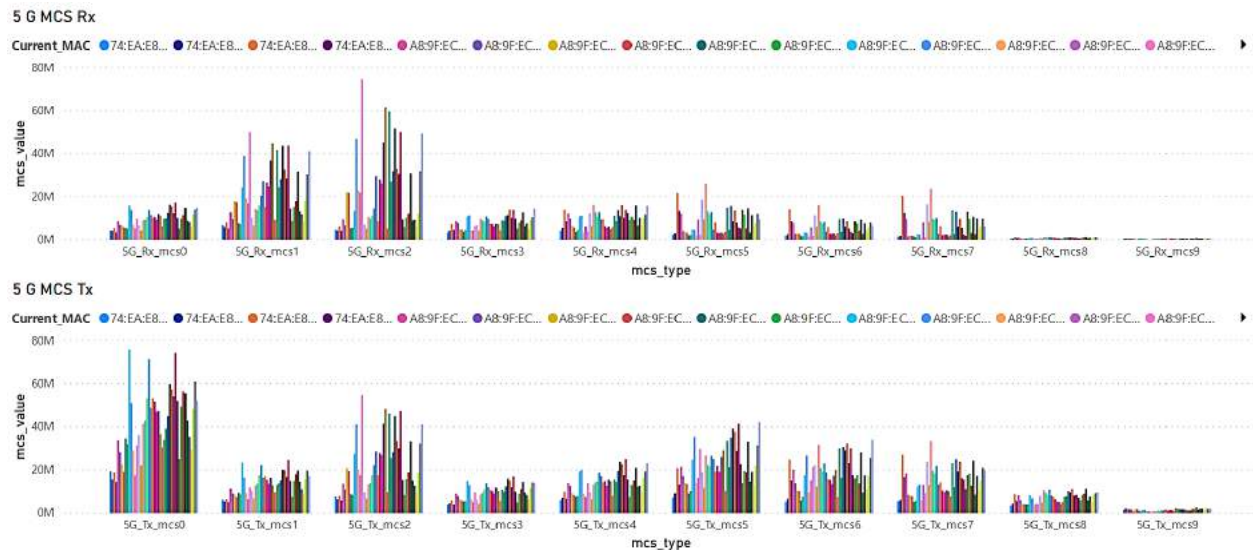
For each of the areas, we also identify the outliers AP. This insight helped in the channel planning, as well as with decisions about turning 2.4GHz radio off entirely in some cases.





**Figure 8 Outliers of Channel Health Statistics**

One of the indications that would provide insight into the way client interact with the AP, in terms of good Signal to Noise Ratio (SNR), Signal Strength Indicator (SSI), and throughput; is the distribution of Data MPDU across the different Modulation and Coding Scheme (MCS) indexes.



**Figure 9 5GHz Radio Data MPDU distribution across MCS**

## 5. Current Results and Findings

### 5.1. Baseline state

First, we needed to establish the percentage of sessions with Bad WCX out of all client sessions.




Figure 10 Percentage of sessions with Bad WCX out of all client sessions – Baseline state

Looking at the breakdown of Bad WCX based on the three types defined above. We can see that moving clients represent the smallest group of bad experience events. Even when looking at outdoor AP exclusively, this non-intuitive fact holds.




Figure 11 Bad WCX breakdown per type – Baseline state

A second insight is that there is no real difference in the Bad WCX due to TCP retries when we look at the percentage of retries compare to all TCP packets in that direction. As expected, the traffic patterns are heavily skewed in favor of downstream. In other words, much more traffic is going from the server towards the client. At first glance, this may be non-intuitive, given the differences in transmit power and receive sensitivity between the client and the AP.




Figure 12 Bad WCX distribution per Radio – Baseline state




Figure 13 Bad WCX distribution per AP Model – Baseline state

## 5.1. Current state

For each of the areas, CommScope has defined several iterations of configuration changes and for the areas that allowed, added coverage by additional outdoor AP or a complete redesign. While COVID-19 shelter-in-place impacted our ability to conduct many of the iterations and significantly reduced the number of observations (client activity, sessions of outdoor activity, etc.). We were able to drive many of the changes based on the insights we gathered by the big data analytics systems.

Using cross-correlation between our defined criteria of Bad WCX, AP SNMP, and telemetry reports. We were able to drive the following improvements on the network. Note the significant improvement in the outdoor AP (OG1600), and the smaller improvement in the average of all models.




Figure 14 Bad WCX distribution per AP Model – Baseline vs. Current




Figure 15 Percentage of sessions with Bad WCX out of all client sessions and improvement




Figure 16 Bad WCX distribution per AP Model – Current state




Figure 17 Percentage of sessions with Bad WCX out of all client sessions and improvement – OG1600s

## 6. Conclusion

Utilizing an automated big data analytics system that collects, correlates, and provides insight into users' experience, even for a nationwide network, is achievable.

Bringing together the right subject matter experts in the areas of data science and Wireless; Operators can use the insights generated by the big data analytic system to identify possible causes of bad user experience. Through recommendations and applying the right actions and implementation, we can eliminate or significantly reduce the number and frequency of these bad user experience events. Paving the way to a more optimized network that can serve more customers and provide significant data offload opportunity for an LTE usage.

Future enhancement of this platform includes more sources of data (e.g., indoor Access Points) and real-time telemetry, which will increase the ability to predict where optimization and intervention are needed. The platform and architecture developed to provide the foundation for a machine learning platform. Such machine learning capability would be able to establish thresholds, identify anomalies, and automate the identification and prediction of user experience. Future enhancement is to close the loop by turning the insights and recommendations provided by the system into an engine that would apply the right configuration changes to the network and continue to optimize the performance.

# Abbreviations

AP	Access Point
bps	bits per second
DHCP	Dynamic Host Configuration Protocol (aka IP allocation)
FEC	forward error correction
GIS	Geographic Information Systems
GPS	Global Positioning System
HFC	Hybrid Fiber-Coax
HHS	Home Hotspot
Hz	hertz
ISBE	International Society of Broadband Experts
KPI	Key Performance Indicator
MAC	Media Access Control
Mbps	Megabits per second
MCS	Modulation and Coding Scheme
MPDU	MAC Protocol Data Unit
MSDU	MAC Service Data Unit
SCTE	Society of Cable Telecommunications Engineers
SMB	Small or Medium Business
SNMP	Simple Network Management Protocol
SNR	Signal to Noise Ratio
SSI	Signal Strength Indicator
SSID	Service Set Identifier (aka. Network Name)
RF	Radio Frequency
RSSI	Receive Signal Strength Indicator
VAP	Virtual AP (aka SSID)
WAG	Wireless Access Gateway

# **An Overview Of Optical Architectures Necessary To Achieve 5G's Key Performance Indicators**

A Technical Paper prepared for SCTE•ISBE by

**Kevin Bourg**

Director, Commercial Technology – 5G Wireless  
Office of the CTO  
Corning Optical Communications  
4200 Corning Place  
Charlotte, NC 28216  
+1 678 464 3200  
kevin.bourg@corning.com

**Sergey Ten**

Technology Strategy Director  
Office of the CTO  
Corning Optical Communications  
tens@corning.com

**Peter Wigley**

Commercial Technology Manager  
Office of the CTO  
Corning Optical Communications  
wigleypg@corning.com



# Table of Contents

<b>Title</b>	<b>Page Number</b>
1. Introduction.....	3
2. Optical architecture considerations to support 5G deployments.....	4
2.1. Network architecture: Point-to-point or point-to-point WDM .....	4
2.2. Network architecture: xWDM and BiDi technology.....	6
2.3. Network architecture: Passive split optical networks (PON).....	8
3. Conclusion .....	10
Abbreviations.....	11

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 The 5G KPIs help define a set of categories which can be applied to use cases for the deployment of 5G services as depicted on this graphic.....	3
Figure 2 High-level overview of the various use cases or categories for 5G services and how they align with different architecture choices (adopted from ITU G.sup.5GP draft) .....	4
Figure 3 - A dedicated point-to-point network architecture employs a separate ODN for access and mobile radio connections.....	5
Figure 4 - A shared point-to-point network architecture has common infrastructure and cable in the feed portion of the ODN for access and mobile transport.....	5
Figure 5 - xWDM converged access networks share infrastructure fiber and connectivity in the ODN extensively.....	7
Figure 6 - Illustration of the various wavelengths utilized within the access network including IEEE and ITU-T PON standard wavelengths and xWDM (illustration courtesy of Mark Hess from Corning Optical Communications) .....	8
Figure 7 - Increases in capacity by PON standards allow for mobile transport and converged residential networks over a common ODN.....	9

## List of Tables

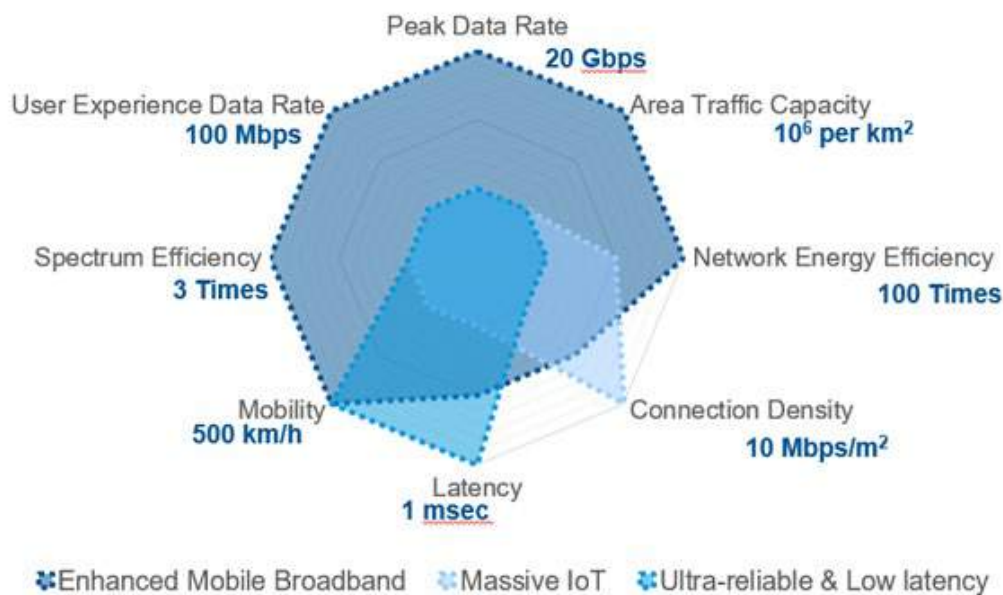
<b>Title</b>	<b>Page Number</b>
Table 1 - Summary of the various attributes of different mobile transport architectures to address 5G network deployments .....	10

# 1. Introduction

Over the past year wireless operators around the world have focused on the deployment of broad 5G coverage. As we are just now beginning to see true 5G devices enter the consumer market, these recent upgrades will support new spectrum options and greater spectral efficiency provided by the 5G standard. The near-term objective is simple: make sure network capacity keeps up with consumer demand.

As we look beyond near-term consumer demand, the 5G standard includes a series of Key Performance Indicators (KPIs) to address a series of use cases beyond today's wireless networks capability. Consider remote driving or e-health use cases where not only bandwidth but also ultra-high reliability communications are required. And possibly the use case supporting rapid forms of transit such as high-speed trains travelling up to 300+ miles per hour. Finally, consider the use case of remote driving cars where one millisecond of latency is critical to avoid a disastrous accident.

Each of the use cases noted and many more can be summarized into three categories: Ultra-reliable and low-latency communications (URLLC), Enhanced mobile broadband (eMBB) and Massive machine type communications (mMTC). Figure 1 shows how the 5G KPIs are assigned to aforementioned categories. To achieve the full potential of one or more categories operators will need to acquire new spectrum and deploy fiber-based radio access network (RAN).



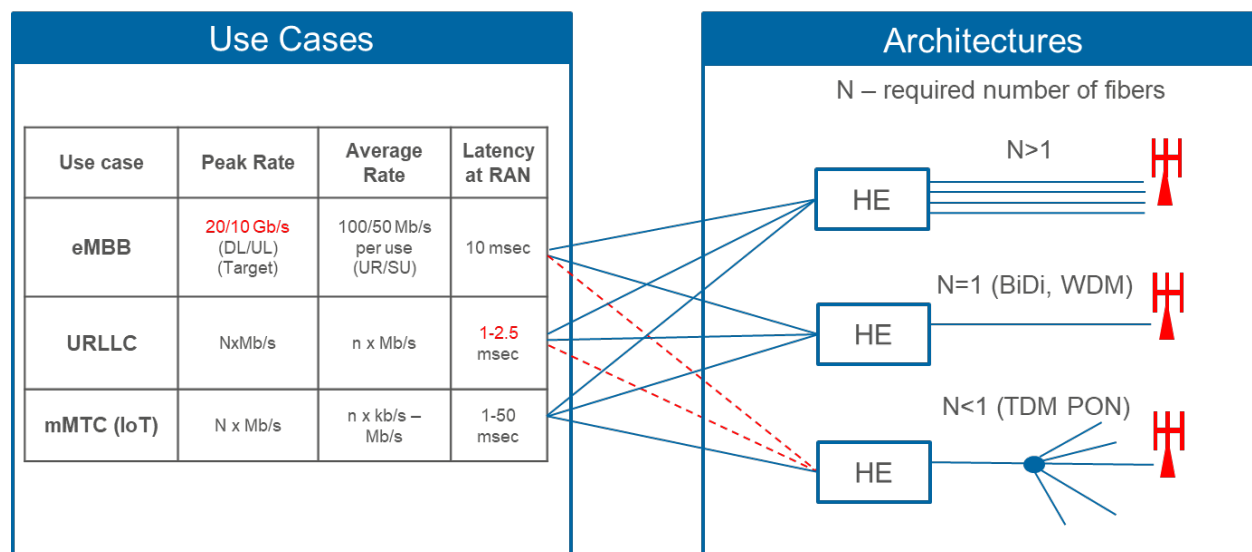
**Figure 1 The 5G KPIs help define a set of categories which can be applied to use cases for the deployment of 5G services as depicted on this graphic**

Existing spectrum deployed today by operators can support 4G subscriber's user experience; however, to address many of the 5G KPIs operators will need to look at new spectrum options such as higher frequency spectrum typically referred to as millimeter wave (mmWave). The mmWave part of the spectrum has large (measured in hundreds of Megahertz(MHz)) available spectral bands to achieve user experience capacity demands. However, mmWave spectrum is more impacted by environmental factors such as rain and snow. In order to overcome those challenges operators must densify their wireless network deploying small cell radios along roadways rather than traditional towers which today have an

Inter Site Distance (ISD) in urban areas on the order of 1,500 feet. The next section will describe a series of architectures that provide the necessary capacity and performance requirements to meet the 5G KPIs. Each of the sections will provide an overview of the architecture, describe some of the salient features of the architecture and provide some guidance on when an architecture should be considered.

## 2. Optical architecture considerations to support 5G deployments

As operators look to densify their wireless networks to support 5G small cell technologies, transport of network capacity from radio head to extensible Radio Access Network (xRAN) locations is predicted to grow at a rate of ~32% per year between 2019 and 2025, according to the Ericsson Mobility Report<sup>1</sup>. (Corning research). This increased demand for 5G transport drives operators to deploy an optical infrastructure. Figure 2 illustrates how the different architectures to be discussed provide support for the various KPIs associated with 5G. The following sections provide guidance on network architecture considerations when densifying a 5G network.

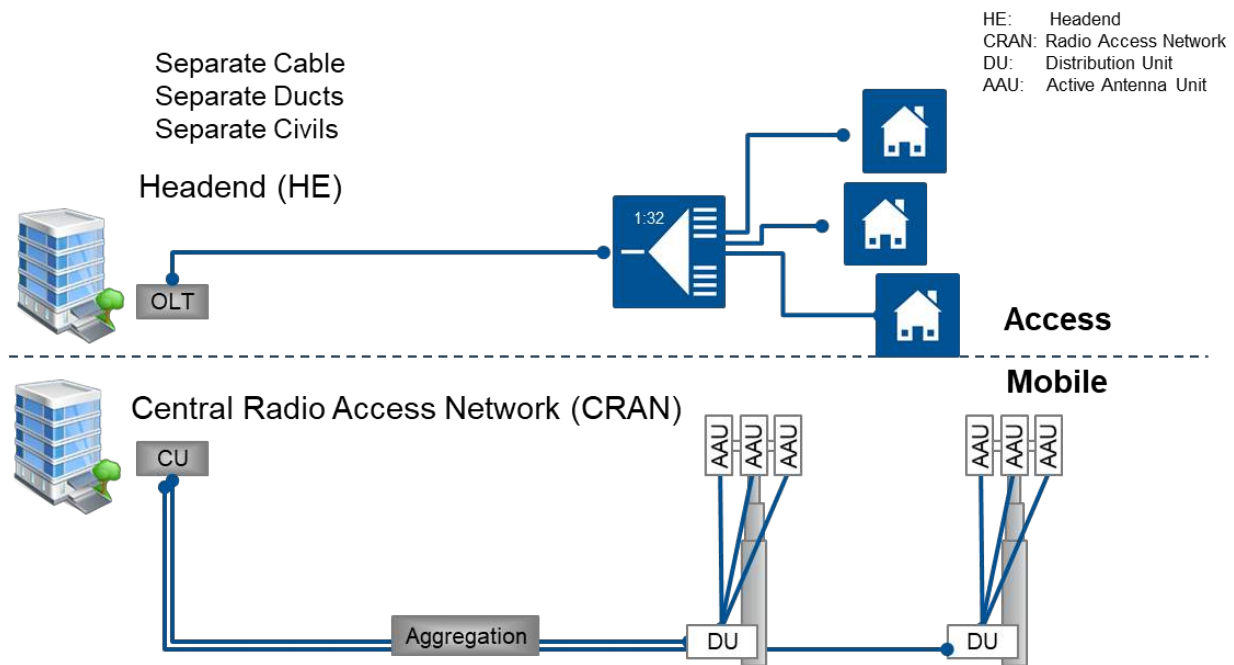


**Figure 2 High-level overview of the various use cases or categories for 5G services and how they align with different architecture choices (adopted from ITU G.sup.5GP draft)**


### 2.1. Network architecture: Point-to-point or point-to-point WDM

The most direct way to provide a high-performance link for mobile transport (from xRAN to the radio) is a fixed, dedicated point-to-point (P2P) architecture. Historically, this is the de facto solution for mobile transport. This architecture requires a dedicated fiber connection from the xRAN location to each drop fiber serving the 5G radio. The optical fibers may reside within a dedicated cable supporting xRAN transport or may reside alongside other optical fibers in the cable supporting other services such as fiber to the home or enterprise-based services. The fiber, cable, connectivity, and infrastructure costs of a dedicated P2P architecture scale as demand for network connectivity grows. Proper up-front planning and over-provisioning of the infrastructure is critical to ensure that long-term growth opportunities will be accommodated with costs.

<sup>1</sup> Ericsson, “Mobile data traffic outlook”. June 2020. <https://www.ericsson.com/en/mobility-report/reports/june-2020/mobile-data-traffic-outlook>



**Figure 3 - A dedicated point-to-point network architecture employs a separate ODN for access and mobile radio connections**



**Figure 4 - A shared point-to-point network architecture has common infrastructure and cable in the feed portion of the ODN for access and mobile transport**

If an operator's existing infrastructure has additional dark fiber within the feeder and distribution cables available for 5G transport or the operator has decided to deploy a new network, a P2P shared architecture is more viable for mobile transport. The low incremental cost and flexibility of dark fiber in a high fiber count cable to support network growth can help manage the uncertainty of future demand (new subscriber

transport, connectivity, mobile densification). Since the incremental cost of fiber is small, sharing cable minimizes incremental cost in the feeder, and sharing duct avoids costly civil work and delays to acquire construction permits if duct capacity is available. Network access points that are already in place might also defer the cost of connectivity if they are adequately sized.

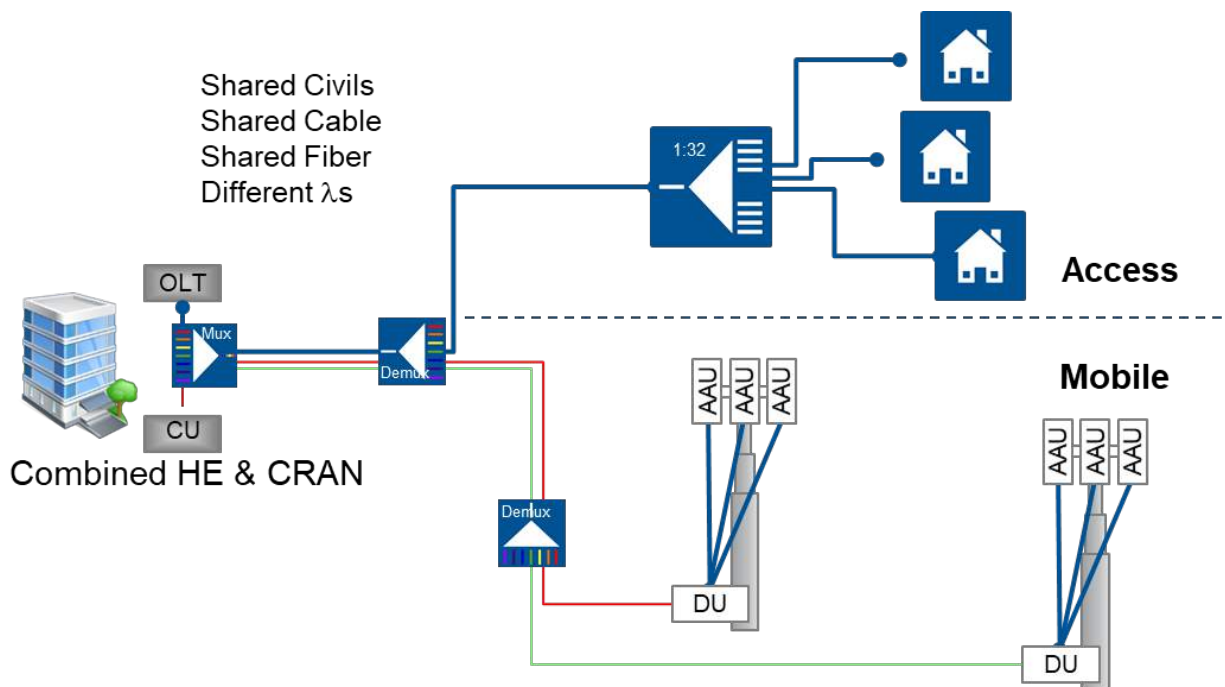
Both dedicated and shared P2P architectures as depicted in Figure 3 and Figure 4 benefit from commoditized, relatively low cost actives, utilizing standardized optics that don't require precise wavelength control or temperature control.

Sharing a cable sheath clearly reduces cable and installation costs in a fiber rich environment, and savings are improved as the number of antennas in the mobile transport links grows. In many markets the availability of dark fiber is low in distribution areas which have experienced significant demand for network infrastructure. These are typically areas where network densification is necessary challenging operators to deploy new optical fiber to support small cell deployment. Hence, operators are very interested in reusing installed fiber in more fiber-lean transport options such as wavelength division multiplexing (WDM) and bi-direction transceivers (BiDi).

## **2.2. Network architecture: xWDM and BiDi technology**

In situations where fiber planning cannot or did not over-provision fiber and cable, or where fiber is exhausted, a dedicated or shared P2P architecture will necessitate new cables and the possibility of significant civil costs if existing duct is not available. An alternative path operators may consider is leveraging WDM technology to extend the capacity carrying capability of exhausted or limited fibers deployed. These technologies enable multiple data channels at different wavelengths to be transmitted through one optical fiber simultaneously. The first and simplest example is the use of Bidirectional optics (BiDi), where two simplex fiber connections for carrying upstream and downstream signals are combined into a single fiber. BiDi technology utilizes a pair of distinct optical signal wavelengths for transmission in downstream and upstream directions. Because the wavelengths are different, upstream and downstream traffic do not conflict. This approach is similar to that used in a Passively split Optical Network (PON), which employ bidirectional transceivers with integrated WDM multiplexers (diplexers) to separate upstream and downstream channels. Since synchronization of upstream and downstream data streams is critical between an adaptive antenna unit (AAU) and distributed unit (DU), BiDi optics minimize the asymmetry in propagation delays since both upstream and downstream signals share a common optical fiber.

BiDi optics are more costly as a result of integrated transceivers with WDM multiplexers, but duplex transmission in a single fiber delays or avoids fiber exhaust, offsetting that cost. BiDi optics also help preserve space in the headend (HE) terminal, providing more shelf space for expansion. BiDi optics are increasingly common, relatively simple to implement, and can be installed in a greenfield or brownfield environment because no additional external optical modules (like WDM multiplexer) required.



**Figure 5 - xWDM converged access networks share infrastructure fiber and connectivity in the ODN extensively**


Both Coarse (CWDM) and Dense (DWDM) xWDM technology are emerging in radio access networks, including unidirectional or BiDi optics. At the transmitter, a xWDM multiplexer combines the output from several optical transmitters for transmission over a single optical fiber. At the receiving end, a demultiplexer separates the combined optical signals and passes each channel to a matched optical receiver. Only one optical fiber per transmission direction is needed between multiple xWDM transceivers, unless BiDi optics are employed, in which case both directions are supported with a single fiber. Since multiplexer and demultiplexer channel ports are specific to a wavelength, transceivers must normally be manually selected to match a specific port's wavelength channel.

Adding an additional mobile transport connection involves connecting a new optical transceiver at the HE or xRAN core to an open optical multiplexer port, and a new optical transceiver to a matched optical demultiplexer port at the radio head. While CWDM technology commonly supports 4, 6, 8 or 12 channels, the much high channel density associated with DWDM technology provides extensive capacity growth for areas with extremely high connection density or where cabling/infrastructure costs/civils are prohibitive.

Most optics in xRAN networks are being deployed at or below 10 Gbps today. More recently, 25 Gbps optics have emerged utilizing CWDM and Local Area Network WDM (LWDM) channel plans, leveraging cost reduced optics resulting from massive deployment of 25 Gbps in hyperscale data center applications. LWDM optics can operate in the O-band where fiber chromatic dispersion (CD) is low, minimizing chromatic dispersion penalty at longer distances, extending the reach of the Optical Distribution Network (ODN).

xWDM technology enables xRAN densification with ODN infrastructure costs that grow as the network grows. WDM multiplexer and demultiplexer technologies increase capital expenses, and xWDM transceiver optics are more costly than standard optics, especially at higher channel density, but once the

initial capital expense is invested, adding additional capacity-becomes a matter of purchasing additional xWDM transceiver channels as densification increases.



**Figure 6 - Illustration of the various wavelengths utilized within the access network including IEEE and ITU-T PON standard wavelengths and xWDM (illustration courtesy of Mark Hess from Corning Optical Communications)**

A xWDM based architecture also offers potential flexibility that cannot be achieved with a fiber rich P2P architecture utilizing standard optics. The origin and destination of traffic in the optical transport layer is normally fixed and inflexible in a dedicated or shared P2P architecture. Since the path in a xWDM architecture is dictated by the wavelength channel selected, multiple mobile antenna sites can be fed from a Macro or xRAN core over a common transport fiber. Extending this further, technology such as reconfigurable add-drop multiplexers (ROADMs) may also be considered to support network topologies with path redundancy enabling the ability to route critical mobile traffic around failure points in the network.

Tunable and auto-tuned optics also offer the potential for simplifying implementation and maintenance through no-touch provisioning. Transport service to a new antenna or other end point can be turned up without having to carry stock and manually select or match transceiver wavelengths to specific WDM ports.

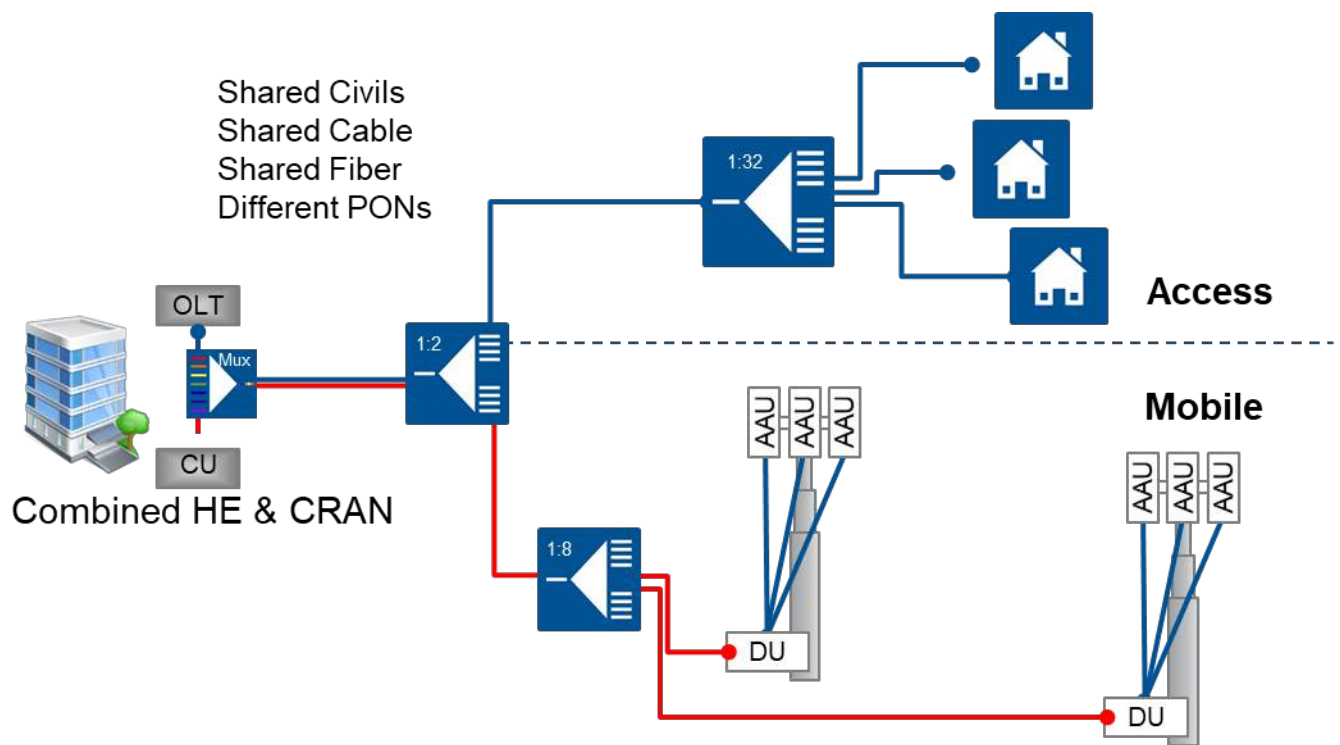
WDM convergence has emerged as way to further extend the capacity and performance of mobile transport networks. Employing a xWDM architecture retains an independent P2P connection from each xRAN terminal at the HE to each radio head, while sharing a common optical fiber in the feeder network.

### 2.3. Network architecture: Passive split optical networks (PON)

Since the early 2000's operators have been deploying fiber to the home networks utilizing passively split optical technology typically referred to as Passive Optical Networks (PON). The attraction of a PON network is the ability to deploy a single optical fiber deep within the network, passively splitting the fiber to support a number (e.g. 32 or 64) of subscribers. In comparison to a point-to-point active ethernet network where there is a dedicated optical fiber from the headend to the subscriber in a 1:1 ratio, within a



PON network the ratio is on the order of 1:32 or 1:64 thereby reducing the electronics needed at the headend and fiber capacity within the feeder network.



**Figure 7 - Increases in capacity by PON standards allow for mobile transport and converged residential networks over a common ODN**

Because a single fiber provides connectivity at the headend to multiple subscribers both ITU-T and IEEE have defined a series of standards to enable a one-to-many Time Division Multiplexing (TDM) (and more recently adding WDM) based communication protocol. This protocol is necessary to ensure upstream traffic from one subscriber device does not conflict with another subscriber. Both standards organizations have continued to evolve their PON standards to increase overall capacity provided over the PON while looking at ways to reduce latency and improve throughput.

Today many operators are deploying symmetrical 10Gbps defined by standards such as IEEE's 10Gbps E-PON or the ITU-T XGS-PON. Other operators are considering much higher capacity PON standards such as NG-PON2 also defined by the ITU-T to provide up to 40Gbps of symmetrical bandwidth over four separate 10Gbps wavelengths. Recently the IEEE released a new standard providing support of either 25Gbps or 50Gbps symmetrical bandwidth. These significant increases in speed provide operators with excellent capacity headroom for the residential subscriber but in many cases are being considered for a converged network supporting both residential subscribers, small enterprises and 5G transport for small cells.

In recent standards development attention has been made to ensure backwards compliance with existing standards or providing multiple wavelengths over a common ODN. Operators could consider utilizing the 10Gbps E-PON standard for residential services while looking at 25G or 50G NG-EPON to provide transport of small cell radios leveraging a common ODN.



As noted above, a PON network utilizes TDM arbitration to ensure each device has a clear channel for upstream communication. This arbitration does highlight a downside of PON networks, increased latency and in some cases increased jitter. In some of the 5G use cases noted above where millisecond latency is critical for communication, a PON based network may not be able to meet the end-to-end service objectives.

The jury is still out as to whether PON can meet the requirements for KPIs that would support all three eMBB, URLLC and mMTC categories in 5G networks. Operators will need to decide on the trade-offs inherent to PON based networks: cost advantages and efficient use of feeder fiber versus increased latency and potential long-term capacity constraints as compared to point-to-point-based architectures. Each operator will need to evaluate these trade-offs and decide what is best for their 5G objectives and network environment.

In Table 1 the authors have summarized a series of different attributes for operators to consider as they look to prepare for the build-out of a 5G transport network. Each of the different architectures have various salient features which will guide an operator on how to proceed based on their current capital constraints and network infrastructure in place.

**Table 1 - Summary of the various attributes of different mobile transport architectures to address 5G network deployments**

	<b>Dedicated PT-PT</b>	<b>BiDi PT-PT</b>	<b>Shared PT-PT</b>	<b>xWDM</b>	<b>PON</b>
<b>Application</b>	Transport	Transport	Transport & FTTH convergence	Transport & FTTH convergence	Transport & FTTH convergence
<b>Fiber count</b>	Highest	High	Highest	Lowest	Low
<b>Adoption</b>	Standard	Emerging	Common	Standard	Under Evaluation
<b>Complexity</b>	Moderate	Moderate	Moderate	Complex	Moderate
<b>Flexibility</b>	Low	Low	Moderate	High	High
<b>Path Redundancy</b>	Low	Low	Low	Capable	Low
<b>Construction cost</b>	Highest	Higher	Lower	Lowest	Lower

### 3. Conclusion

The growing demand by consumers on the mobile network is predicted to increase at a rate of 30% per year. This demand alone will push existing wireless technology and architectures to the brink of capacity. The deployment of 5G wireless networks does provide operators an opportunity to gain additional spectral efficiency; however, this alone will not be enough. Further, emerging use cases are pushing additional capacity, latency and connectivity requirements, demanding new spectrum and increasingly dense radio networks.

This paper provided an overview of essential architectural approaches, describing deployment issues to consider based on a carrier's use cases. The authors are also aware that combinations of these architectures as well as other technologies are worthy of consideration. These may include the possible use of Cable Labs Data-over-Cable Service Interface Specifications (DOCSIS®) or Integrated Access-

Backhaul (IAB) technology to provide this transport from the radio site to the xRAN location (or intermediate location). Although these are acceptable technologies, we challenge each operator to look over the 10-year horizon to ensure the infrastructure dollars spent today are capable of meeting the demands of the longer-term network demands.

## Abbreviations

AAU	Active Antenna Unit
BiDi	Bi-direction Transceiver
CD	Chromatic Dispersion
CRAN	Centralized Radio Access Network
CU	Central Unit
CWDM	Coarse Wave Division Multiplexing
DU	Distribution Unit
DWDM	Dense Wave Division Multiplexing
eMBB	Enhanced Mobile Broadband
HE	Headend
IAB	Integrated Access-Backhaul
ISD	Inter Site Distance
KPI	Key Performance Indicator
LWDM	Local Area Network Wave Division Multiplexing
MHz	Megahertz
mMTC	Massive Machine Type Communications
mmWave	Millimeter Wave
ODN	Optical Distribution Network
P2P	Point to Point
PON	Passive Optical Network
RAN	Radio Access Network
ROADMs	Reconfigurable Add-Drop Multiplexers
TDM	Time Division Multiplexing
URLLC	Ultra-reliable And Low Latency Communications
WDM	Wavelength Division Multiplexing
xRAN	Extensible Radio Access Network

# **Delivering Cloud-Native Operations with Edge Compute Enabled DAA**

## **Implementing a Kubernetes Distributed Edge**

A Technical Paper prepared for SCTE•ISBE by

### **Marco Naveda**

Sr Director, Network Architecture, Office of the CTO  
Ciena  
5050 Innovation Drive Kanata, ON K2K3K1 Canada  
613-670-2730  
mnaveda@ciena.com

### **Dmitri Fedorov**

Software Architect, Office of the CTO  
Ciena  
5050 Innovation Drive Kanata, ON K2K3K1 Canada  
(613) 670-2757  
dfedorov@ciena.com

### **Raghu Ranganathan**

Principal, Advanced Architecture, Office of the CTO  
Ciena  
7035 Ridge Rd, Hanover, MD  
713-662-9999  
rraghu@ciena.com

# Table of Contents

Title	Page Number
1. Introduction.....	4
2. Drivers for Access Network Modernization .....	4
3. Cloud Native and Edge Computing Architectures .....	6
3.1. Distributed Edge Computing.....	6
3.1.1. Properties of Distributed Edge Computing .....	8
3.2. Cloud Tiering Reference Model .....	8
3.2.1. 3-tier Cloud Model.....	8
3.2.2. Elements of the Edge Tier.....	9
4. Cloud Native Technologies.....	10
4.1. Cloud Native Applications .....	10
4.2. Virtual and Cloud Native Network Functions .....	12
4.3. Container Orchestration.....	14
4.3.1. Kubernetes .....	14
4.3.2. Kubernetes Distributions .....	16
4.3.3. Vanilla Kubernetes Distribution.....	16
4.3.4. Kubernetes with Value-add Capabilities .....	16
4.3.5. Kubernetes as a Service .....	17
4.3.6. Lightweight Kubernetes for Edge Cloud and IoT.....	17
4.3.7. Kubernetes for Distributed Node Clusters.....	17
5. Building out a Distributed Edge Cloud .....	18
5.1. Cloud vs Edge Orchestration .....	18
5.2. Network Functions and Runtimes .....	19
5.3. Automated Operations & Business Network Intent .....	20
5.4. Application Repositories .....	20
5.5. Undercloud Architecture .....	21
5.5.1. Undercloud Control Plane .....	22
5.5.2. Undercloud Resources & Edge Optimized Runtime.....	24
5.6. Open-Source Building Blocks .....	24
5.6.1. LF Edge.....	24
5.6.2. Cloud Native Computing Foundation .....	26
5.7. Commercial Cloud Platforms for Edge Computing.....	26
5.7.1. AWS IoT Greengrass .....	26
5.7.2. AWS Outpost.....	27
5.7.3. Azure IoT Edge.....	27
5.7.4. Azure Stack Hub.....	27
5.7.5. Google Cloud Anthos .....	27
6. Conclusion .....	27
Abbreviations.....	27
References.....	29

## List of Figures

Title	Page Number
Figure 1 - MSO transition to DAA and fiber deep.....	5
Figure 2 - Edge Application Segmentation .....	6
Figure 3 - Edge Locations.....	7
Figure 4 - Three-tier Cloud System .....	8

Figure 5 - Edge Cloud Network Reference Model .....	9
Figure 6 - Cloud Native Constructs .....	11
Figure 7 - VNF vs CNF .....	13
Figure 8 - Kubernetes Components .....	15
Figure 9 - Kubernetes Distribution Models .....	16
Figure 10 - DAA Orchestration Framework.....	18
Figure 11 – Intent-Driven Undercloud .....	22
Figure 12 - Homogeneous Resource Scheduler.....	23
Figure 13 - Heterogeneous Resource Scheduler.....	23
Figure 14 - Akraino Software Stack.....	25

# 1. Introduction

The networking software industry is experiencing an accelerating technology shift from centralized data center application delivery models to a distributed edge computing paradigm. In this new computing paradigm, cloud infrastructure and services are delivered from multiple distinct and geographically distributed locations in closer proximity to the end user or data source. The drivers for this shift are the emergence of advanced Enterprise use cases & applications, network infrastructure convergence and operator digital transformation initiatives. These applications require deterministic latency and ultra-fast response times, distributed processing of large volumes of data near the source and flexible scalability across network and computing ecosystems. Included in this transition, the convergence between wireless & wireline, combined with increased hub-site network capacity in cable Distributed Access Architectures (DAA) and the drive to centralized mobile Radio Access Networks (cRAN), are making distributed edge computing more feasible for operators.

This computing paradigm is enabled by cloud-native principles and application containerization & orchestration technologies that promise new operational efficiencies and agility to develop innovative services. We see this trend across Enterprises and network operators that want to accelerate software delivery while maintaining a consistent quality of experience from applications & data to devices and end-users alike, independent of location. 5G and high-speed fixed broadband connectivity services are acting as catalysts to enhance the underlying network performance and agility using a cloud-native services-based architectures. This in turn poses a challenge of operationalizing an applications-first approach in the operator's network to facilitate open, modular, and portable multi-vendor networking software across their distributed edge compute infrastructure platforms.

Advanced cloud-based management & orchestration systems, however, were not designed for distributed edge computing. These systems were optimized for very large compute environments where server clusters are co-located and mesh inter-connected by an over-provisioned data center fabric. In edge centric architectures, compute & network nodes are geographically distributed and inter-connected by multi-layer access and aggregation networks with varying degrees of capacity, latency, and flexibility. This creates the opportunity to adapt and optimize the use of containerization and cloud orchestration technologies to meet the needs of embedded real-time network functions and edge business applications.

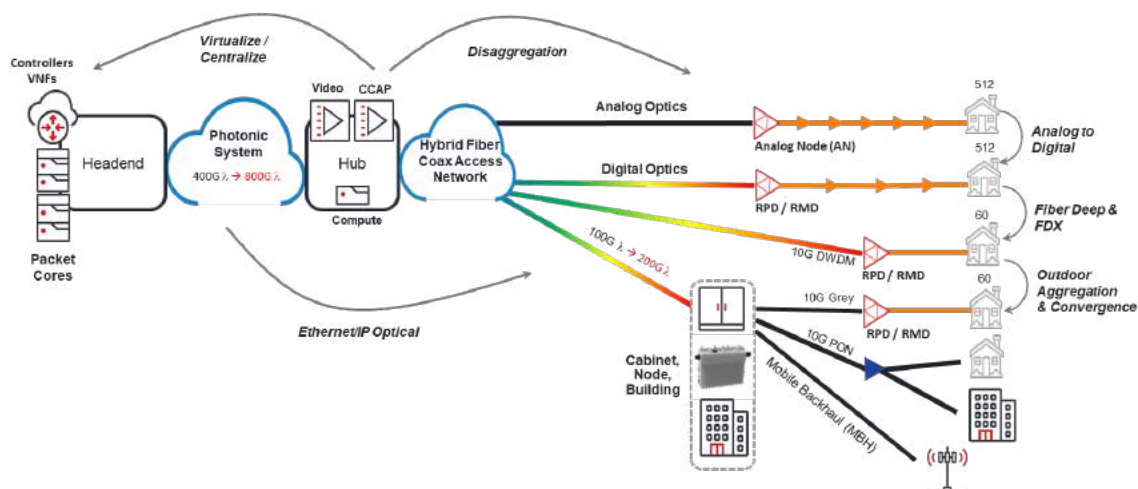
This paper will delve into emerging edge computing infrastructure architectures, available open-source software to enable distributed edge computing, and key technical considerations to accelerate adoption of new software-based operational methods. This paper evaluates the use of open-source projects from the Linux Foundation Edge (LF Edge) and Cloud Native Computing Foundation (CNCF), such as Akraino Edge Stack and Kubernetes, to build and operationalize distributed edge computing networks. This paper will also explore how hyperscale cloud platforms are addressing this technical challenge and how cable MSOs can leverage a rich technology ecosystem to architect open edge technology systems, implement software defined network operations, and monetize new edge-based business services.

## 2. Drivers for Access Network Modernization

In the ever growing quest for delivering higher reliability, higher speed connectivity services to residential and business customers, MSOs are overhauling their traditional Hybrid Fiber Coax (HFC) networks along three key dimensions: (1) disaggregation and distribution of vertically integrated, proprietary functions of the CMTS previously located at headend or hub locations; (2) deeper fiber-based Ethernet/IP connectivity to remote access infrastructure nodes where MAC and PHY layer processing occurs; (3) convergence of service functions & infrastructure for both wireline and wireless services.

Functional disaggregation based on the DOCSIS DAA specifications allow for splitting of MAC & PHY layers from the CMTS, driving either a Remote MACPHY Device (RMD) or a Remote PHY Device (RPD) deeper into the access plant (Levensalor & Stuart, 2020). DAA extends the digital processing of the headend and hub domains out to fiber nodes, pushing intelligence closer to the end user and offering the opportunity to leverage generalized compute platforms, NFV and SDN. This in turn facilitates flexible deployment, optimized infrastructure and automated operations which would otherwise be increasingly complex due to the distributed nature of this architecture (Evolution to Distributed Access Architectures, n.d.). This is illustrated in **Error! Reference source not found..** Additionally, since these locations are typically constrained in space and power, there is a need for light-weight, high-performance compute resources and efficient virtualization technologies such as Linux containers to maximize use of CPU and network resources.

Pushing fiber and Ethernet/IP connectivity deeper in the access plant has the effect of increasing the available access network capacity to a smaller number of homes or businesses in a service group. This in turn allows operators to boost service bandwidth and reliability, while guaranteeing SLAs via increased levels of visibility and control in the physical network underlay. A unified and automated physical & virtual network underlay is a critical component of orchestrating edge compute & storage infrastructure in support of dynamic infrastructure and overlay end-user applications and services (The Converged Interconnect Network, 2020).



**Figure 1 - MSO transition to DAA and fiber deep**

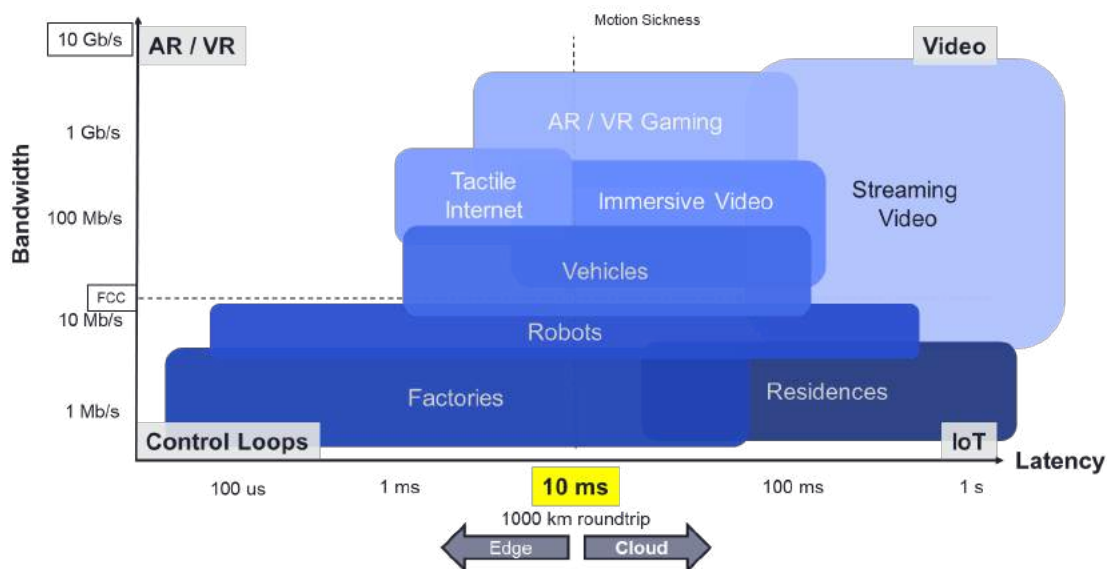
Finally, Packet Core functions delivered via a virtualized CCAP such as user management and authorization are non-data/user plane software-centric capabilities that can be delivered out of a centralized location leveraging the economies of scale offered by cloud-based IT platforms. These functions can be placed in headend or regional data center locations but would necessitate a unified approach to manage the end-to-end lifecycle of open, modular, multi-vendor software components that ran as a monolithic application in a single-vendor CMTS system. This comes with the added benefit of enabling the convergence of service delivery functions across different access medium, such as FTTH and 5G RAN.

### 3. Cloud Native and Edge Computing Architectures

#### 3.1. Distributed Edge Computing

As the cost of commodity computing hardware continues to decline and smart devices and sensors shrink in size, it becomes economically feasible to build connected infrastructure that is continuously monitored and optimized using data-driven insights and intelligent automation systems. In addition, 5G technology promises to enable the interconnection of tens of billions of devices, sensors, and things, so we are going to see an exponential increase in endpoints coming online and generating massive amounts of data that need to be processed closer to its source.

When we overlay this technology landscape with new kinds of business oriented real-time, low-latency applications, like self-driving cars, robotic manufacturing, industrial process automation, and augmented/mixed reality, it becomes critical to segment the application space based on bandwidth and response time requirements. See Figure 2 - Edge Application Segmentation. This high-level segmentation provides a framework to think about applications in terms of its functional components, communications requirements and where these components need to be located to meet the end user experience. Since most public cloud regions are within 60-100 msec of high-density population centers, it becomes clear that a good subset of these high intensity applications is not feasible with today's centralized cloud model.



**Figure 2 - Edge Application Segmentation**

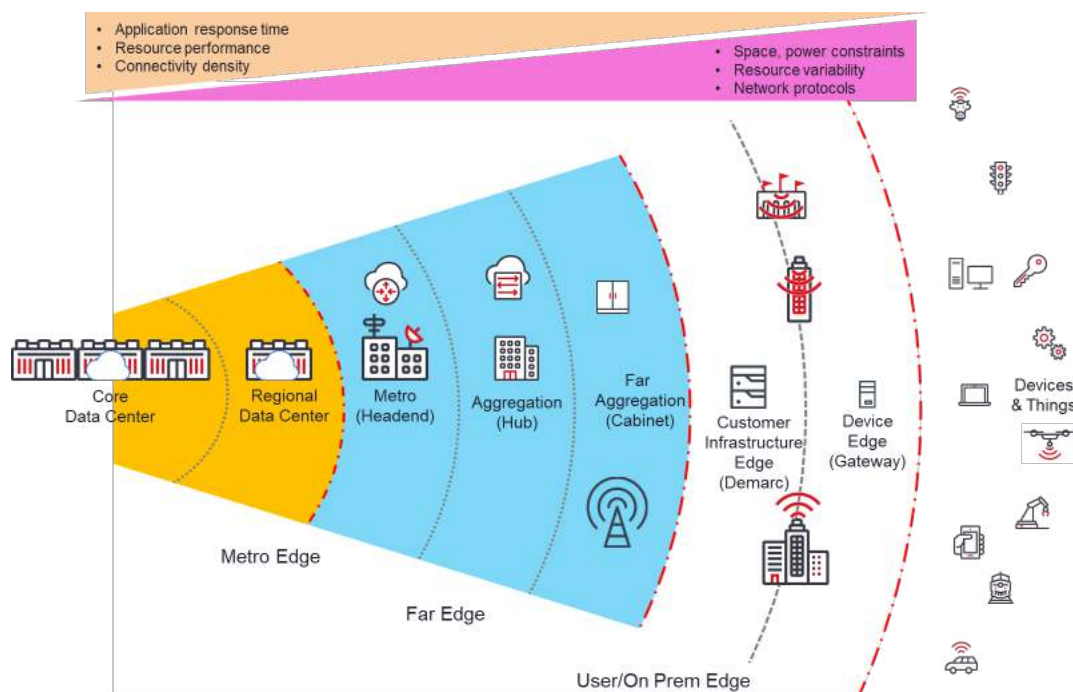
Centralized cloud computing has been a huge success by all measures, and businesses continue to migrate both generic IT and specialized workloads to public clouds. These businesses include Communication Service Providers (CSP) and network operators who are embracing cloud-native technologies to virtualize their networks, modernize operations, accelerate the pace of service innovation, and dramatically reduce costs. The resulting cloud workloads range from business-oriented IT applications such as planning, order management and service assurance, to network-centric software applications such as network service orchestration, virtualized routing and other CPU-based data / user plane functions. These workloads will be distributed across heterogeneous central and edge-based data centers that require a unified approach for software management and operations.



On the technology front, one of the corner-stone technologies of cloud computing – server virtualization or virtual machines – is increasingly being replaced by Linux container technology which has three compelling advantages: rapid workload deployment, better CPU utilization and lighter-weight resource footprint. This container technology supported by tools like Docker have the added benefit of simplifying application life-cycle, from creation and packaging, to testing and delivery across any infrastructure running a Linux OS.

These trends are helping drive an industry shift to create a more distributed and infrastructure optimized computing model, which moves cloud-style consumption of compute, storage, and network resources closer to the end-user or data source. We call this new computing paradigm Distributed Edge Computing (DEC), whereby a software-centric approach facilitates dynamic placement of application components across a heterogeneous environment of connected compute & storage resources, while abstracting the complexities of operating these resources from the application developer. These resources may reside on a smart sensor powered by an ARM processor, an IoT gateway running on a network appliance or a data center located within a hub location of an MSO network. See Figure 3 - Edge Locations.

This DEC approach is also fundamental in the move towards virtualization and distribution of CMTS functions that must be dynamically, but intelligently placed across Hub and access infrastructure with the right resources to meet performance requirements, such as Low Latency DOCSIS (LLD) specifications. LLD targets 1ms queuing and overall less than 10 ms round trip time in the access network (Whie, Sundaresan, & Briscoe, 2019).



**Figure 3 - Edge Locations**

### 3.1.1. Properties of Distributed Edge Computing


There are in fact multiple names and definitions of this style computing, so we propose that there are three key properties shared across the spectrum of edge-centric computing architectures.

- *Location*: physical location of a compute node in a distributed environment defines communications latency, data sovereignty and ability to perform specialized hardware-assisted processing. This covers a spectrum of compute capabilities from a sensor/device in a mining field to customer premise and operator network infrastructure hosting network and application services within a metropolitan area.
- *Heterogeneous cloud*: distributed edge computing is an optimization of central cloud computing, and as a result it inherits properties such as shared resource pooling, elasticity, and application agnostic infrastructure. This is often overlooked but it emphasizes the need to support a holistic and cloud-native approach to containerized application delivery and multi-location data processing according to resource constraints, performance & quality of experience.
- *Network diversity*: as application components and supporting resources become more distributed, the underlay communications network becomes more diverse, and increasingly critical, in terms of protocols, technology domains, application awareness and transport flexibility. This is in contrast with network applications that run within a server cluster in the same core data center.

## 3.2. Cloud Tiering Reference Model

### 3.2.1. 3-tier Cloud Model

From a software application standpoint, edge computing infrastructure fits into a 3-tier system, where the edge tier is located between the hyperscale cloud and devices to perform specialized functions. This 3-tier system limits the extent to which device level applications need to communicate with a centralized cloud for active data storage, processing and other services, as illustrated in **Error! Reference source not found.**



**Figure 4 - Three-tier Cloud System**

When this edge tier is missing, which by and large is the norm today, the system becomes a traditional cloud computing environment whereby all the application intelligence is centralized and possibly assisted by smart devices. From an application perspective, the edge tier can play a dual role to deliver network services and application services. In a 5G RAN, the edge can host radio signal processing functions and user plane functions (UPF) for local traffic breakout and termination at the application layer. Another key

role is to bring compute resources closer to the device tier to improve application response time, reduce unnecessary data transfers to the cloud tier and enhance communications security.

### 3.2.2. Elements of the Edge Tier

The edge tier's functional role also varies according to the capabilities and ownership of the underlying communications network. The edge compute tier sits between the enterprise LAN and operator's metro core to leverage last mile networks and enable local traffic breakout functions for general traffic off-load, time-sensitive end-user applications as well as network-centric protocol processing supporting higher-level applications. As a result, the edge tier can be further decomposed into the Device Edge and the Infrastructure Edge as illustrated in **Error! Reference source not found..**

The purpose of the Device Edge is to host a Device Edge Cloud or sometimes referred to as On-Prem Edge Cloud where enterprise owned devices can benefit from proximity and security of an on-net cloud environment for local application processing and storage serving the enterprise environment. This is particularly useful where enterprises need to retain control of their network traffic and where their data is processed and stored.

The Infrastructure Edge is located on the operator side of the last mile network and typically hosts an Edge Cloud environment for Telco-centric workloads, or what is commonly called a Telco Cloud. These environments are owned and operated by last mile network operators and are becoming increasingly attractive to offer internal network services, shared infrastructure wholesale services and enterprise business services in collaboration with hyperscale cloud providers. The infrastructure edge is where we see the potential for MSOs to differentiate their network services by creating access on-ramps onto an edge cloud for running gaming, AR/VR, Smart City IoT and other real-time analytics intensive applications near the data source.

The Edge Cloud represents not just an architectural choice, but also a system that encompasses storage and compute assets located at the edge of the network, and interconnected by a scalable, application-aware network that can sense and adapt to changing needs, securely and in real-time.

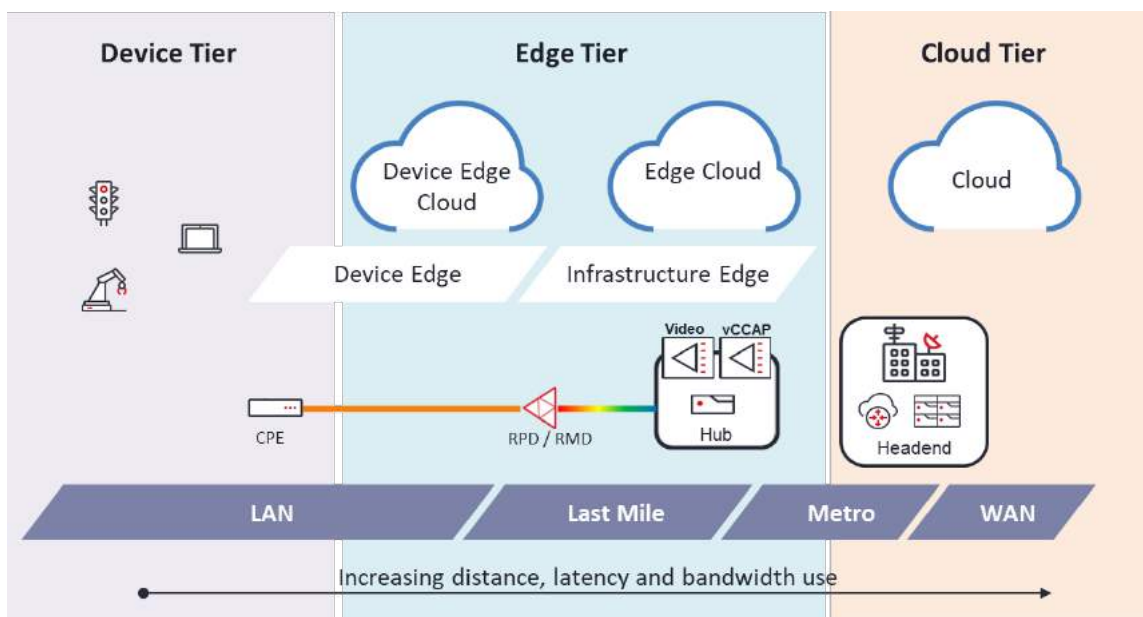


Figure 5 - Edge Cloud Network Reference Model

From an infrastructure perspective, the same cloud networking model can be applied to an MSO network, with the device tier defining the subscriber's domain and the edge tier encompassing the outside plant passive equipment, RPD/RMD nodes as well as hub locations hosting vCCAP systems. The hub and headend locations are equipped with data center infrastructure such as server clusters and cloud-native platforms to deliver the compute and networking environment for virtualized packet core and video functions. These are considered edge data centers due to proximity to the subscriber, footprint requirements and variety of network functions requiring specialized hardware depending on DAA design choices.

On the other hand, RPD/RMD nodes are not multi-server clusters, but they can be considered generalized compute node extensions of the infrastructure at the hub or edge data center, and therefore can participate in the end-to-end orchestration of resources allocated to user plane network functions in the last mile. Due to the nature of MAC/PHY processing functions, such as FEC and MAC scheduling, these nodes can benefit from having specialized accelerators like GPU, TPU, FPGA and smart NICs with very efficient techniques for pooling and scheduling micro-workloads in a distributed fashion. These edge nodes are typically in space and power constrained locations but can also be deployed in edge data center environments where a remote CMTS is desired due to population density.

One key property of these edge nodes and small edge data centers is their highly dispersed physical locations, which require high degrees of autonomous operation and resilience. Given the underlying Ethernet/IP fabric, there is an opportunity to turn a mesh of edge nodes and data centers into one pool of resources for dynamic orchestration of virtualized edge functions, when and where they are needed, thereby reducing the total cost of ownership for the operator.

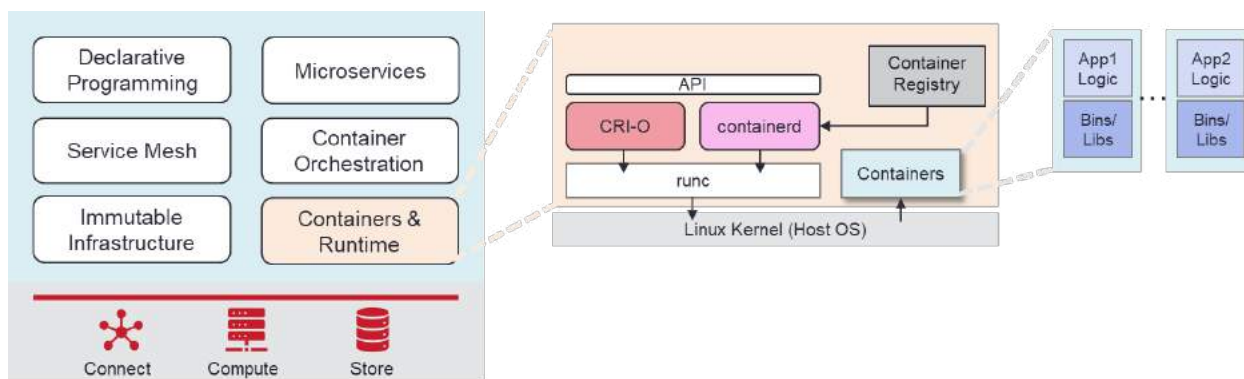
## **4. Cloud Native Technologies**

### **4.1. Cloud Native Applications**

With the development and proliferation of hyperscale cloud platforms from Amazon, Microsoft, Google and others, a new kind of software development and delivery paradigm has emerged called *cloud native* applications. This approach is based on the principle of decomposing an application into a set of microservices that can be developed and deployed independently to accelerate & optimize the DevOps life cycle of software systems. These microservices are packaged into light-weight containers which are scheduled to run on compute nodes by a container orchestrator. There are many advantages of containers

vs virtual machines, but the principal ones are portability, low resource usage, dynamic horizontal scalability, and fast restarts.

*Cloud native* applications are built to run and scale in public, private, and hybrid clouds and they use the following constructs to deliver on the promise of a developer-centric approach to enable cloud scale agility, scalability and flexibility of software systems. See Figure 6 - Cloud Native Constructs.



**Figure 6 - Cloud Native Constructs**

Containerization is an operating system virtualization paradigm in which the kernel supports multiple isolated user space instances or namespaces. From an application perspective, a container is an executable binary packaged with its lib dependencies and intended for execution in these private namespaces with resource constraints such as CPU, memory and storage. The lifecycle of a container is managed by what is commonly called a container runtime. There are several container runtime implementations, each with their own approach at managing the end-to-end lifecycle of a container. In Linux, the execution phase of a container is generally performed by runc ([github.com/opencontainers/runc](https://github.com/opencontainers/runc), n.d.), an Open Container Initiative compliant implementation, but alternatives such as *rkt*, pronounced “Rocket” ([coreos.com/rkt/](https://coreos.com/rkt/), n.d.), are also available. The configuration and image management are performed by applications such as docker, containerd and cri-o which interact with runc. In an effort to simplify the integration with the different container runtime flavours, the Kubernetes Container Runtime Interface (CRI) offers an abstraction layer to interact with the underlying container runtime.

Using containers instead of virtual machines increases CPU utilization and significantly reduces disk space requirements. This is because containers running on the same host share the operating system (OS) while virtual machines have their own OS, providing complete isolation between apps. This property makes containers a more attractive choice for running software-based network functions and applications at the Edge Tier’s compute, power and space constrained hardware.

Service mesh is a dedicated infrastructure layer for service-to-service communication. This infrastructure includes not only network connections between containers and services they form, but also a means of discovering services. This layer makes possible direct communications between containers under policy control (therefore the term *mesh*).

Microservices are loosely coupled fine-grained services with lightweight communication protocols. This architectural style was developed as an alternative to large, tightly coupled services. It brings not just

modularity and scalability vital for applications at the edge cloud, but also supports incremental integration with legacy systems and distributed, parallel development of software (Namiot & Sneppe, 2014).

*Immutable infrastructure* is the concept of never requiring server infrastructure to be modified to support new requirements, but rather new servers are built to replace the old ones. This approach reduces operational complexity by eliminating the need to deal with differently upgraded systems and makes possible quick and fully automated recovery from faulty software that was rolled out to customers. Also, when implemented as the foundation for container images and their workloads, immutable infrastructure removes from consideration the difference between development, test, and production environments.

*Declarative Application Programming Interface (API)* is a design style that avoids specifying *how* to perform described functions, instead, it describes *what* needs to be done. This style allows understanding and consuming of services without knowledge of the services implementation. This reduces integration complexity and promotes service modularity and scalability, simplifying operations.

Cloud native applications consist of loosely coupled, resilient, manageable, and observable container-based microservices. DevOps teams use automation to make high-impact changes frequently and predictably with minimal effort within large scale data centers. Applications designed for the Edge Tier infrastructure use the same cloud-native principles, but must take into account the resource constraints and location context characteristics mentioned in Properties of Distributed Edge Computing.

*Edge Native* applications are impractical or undesirable to run in centralized data centers at public, private, and hybrid clouds. These applications are developed with proximity and specialized resources in mind as well as different security, compliance and networking requirements due to location. Edge-native applications use the infrastructure edge to provide large-scale data ingest, data reduction, real-time decision support, bandwidth savings or to retain sovereignty over critical data.

In the spectrum between Edge Native and Cloud Native Applications, Edge Enhanced is a set of applications that can operate in a centralized data center, but would gain performance, typically in terms of latency, or functionality advantages when operated using edge computing. These Edge Enhanced applications may be adapted from existing cloud native applications or may require no changes if the edge cloud environment is abstracted by the container runtime engine and associated container networking facilities.

## **4.2. Virtual and Cloud Native Network Functions**

Cloud-native principles can be applied to enterprise and consumer-oriented applications, as well as communications and networking software that we refer to as Telco or MSO workloads. These workloads can be broadly classified according to the function and services offered in the network operator's stack:

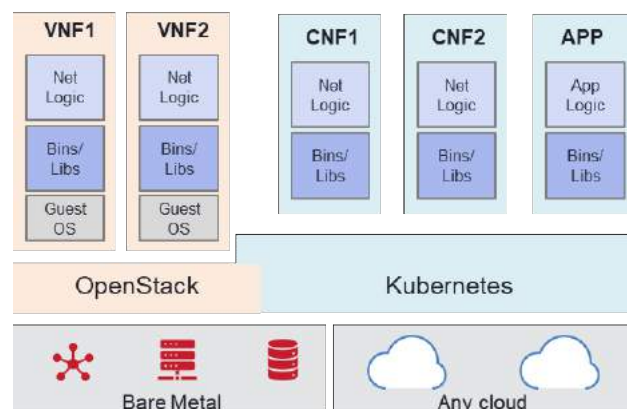
- *Management*: Business Support System (BSS) and Operations Support System (OSS) are software solutions that help operators manage the customer experience, network offerings and network operations for planning, engineering, ordering, activating, and assuring communications services. This type of application is non-real-time and does not interact with the user / data plane & network traffic directly.
- *Control*: network management, signaling & control plane software solutions provide network device inventory, discovery, configuration, performance, fault and resource management capabilities in support of user / data plane services, but are not directly involved in processing the subscriber's / customer's traffic. This includes systems such as user authentication & session management, service policy enforcement, and DHCP services for IP address assignment. These

network functions are considered management and control functions that can be deployed as virtualized or containerized applications without significant differences in network user plane performance.

- *User plane*: these are multi-layer network traffic processing functions that operate at the network layer and below for the purposes of packet inspection, encapsulation, transformation, QoS treatment, forwarding, and filtering among other functions. These are the most-real-time intensive network functions that can operate in a subscriber's cable modem or uCPE, RPD/RMD node or headend location. This is the type of networking software stack that is impacted the most when migrated to a virtualized or containerized environment due to its performance and reliability requirements, but more importantly the life-cycle management processes and techniques used by the vendor and network operator community.

Network Function Virtualization (NFV) has been in production networks for several years now, including data plane functions such as virtual routing (vRouter), virtual firewall (vFW) and virtual Broadband Network Gateway (vBNG). These software-based data plane functions, or Virtual Network Functions (VNF), evolved from their dedicated physical appliance counterparts. When a network function is implemented as a software stack that runs in a virtualized compute environment, as a Virtual Machine (VM), it is referred to as “virtual” or “virtualized network function” (VNF). In many cases, several VNFs will operate in an edge cloud as part of a network service chain that provides a composite service to the end customer.

Network functions can run inside a virtual machine or a container. When a network function is built and deployed as a cloud-native application it is referred to as “cloud native network function” (CNF) or “containerized network function”. See Figure 7 - VNF vs CNF. This means that the software is distributed as a container image and deployed, managed, and orchestrated by tools like Docker and Kubernetes. Both VNFs or CNFs support lifecycle operations that enable frequent and automatic deployment and updates of the software; this is fundamentally different from traditional processes where network operations update network elements on a controlled and infrequent basis.



**Figure 7 - VNF vs CNF**

It is important to note that even though the software capabilities of a VNF may be containerized, the VNF itself is not orchestrated as a containerized application because the delivery mechanism for deployment is a VM image. In other words, containers inside the VM are not exposed to an external container orchestration system. Container orchestration or lifecycle management inside a VNF is typically handcrafted by the VNF vendor using tools like Docker compose and therefore container resource management is limited to the resources allocated to the VNF instance at deployment time. The



implication is that mechanisms for the operator to deploy, scale, monitor and heal these software-based network functions are very different and they operate at different abstraction levels, VM-based VNFs running on virtualization systems such as OpenStack or VMware vs container-based CNFs running on a container runtime such as Docker and orchestrated by Kubernetes.

These differences are critically important at the infrastructure edge where real-time user plane functions need to make optimal use of resources and require direct access to acceleration hardware, where containers can be much more efficient. Another key consideration for edge network functions is the location of the control plane responsible for lifecycle operations across a distributed set of compute elements hosting VNFs or CNFs. In the case of containerized functions, this control plane is increasingly based on Kubernetes which is based on a collocated server cluster concept. More on this in the next section, but the reliability and make-up of underlying network infrastructure will impact the control plane design.

### **4.3. Container Orchestration**

#### **4.3.1. Kubernetes**

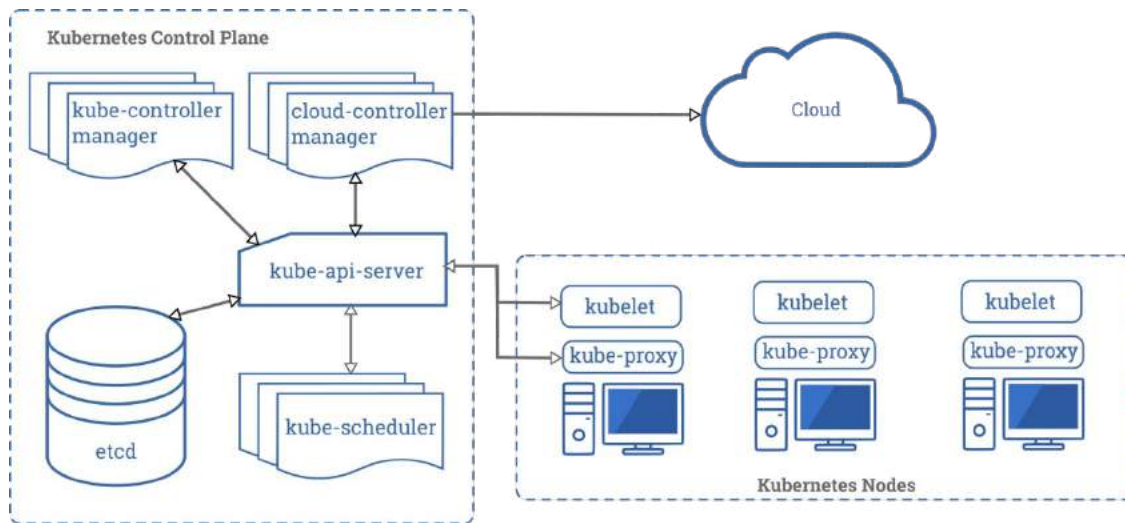
Kubernetes is one of the most widely used container orchestration and management platforms in the Enterprise IT industry (Alusha, 2019). The rise of Kubernetes (K8s) has enabled cloud providers to offer managed K8s services that allow enterprises to create a hybrid / multi-cloud environment for their applications. K8s is an open-source system for automating deployment, scaling, and management of containerized applications (Kubernetes, automated container deployment, scaling, and management, 2020). K8s simplifies the deployment of scalable, distributed applications by managing the lifecycle of containers, including scheduling, load balancing and distribution across different server nodes.

Kubernetes was designed for large scale cloud environments, and it works well out-of-box only when the infrastructure edge consists of one or more Kubernetes clusters and their master nodes have a fast and reliable connection to worker nodes. The Kubernetes master node (or control plane) is relatively heavy, and while its worker nodes are less resource demanding, they are still not lightweight at all. While specific sizes vary depending on its distribution and version, Kubernetes best practices for running large clusters (Kubernetes Best Practices, n.d.) recommends at least 4 GiB of RAM and a single core Intel Xeon CPU for a master node that controls up to 5 worker nodes. When the number of nodes grows, so grows the RAM and CPU requirements for the master node, getting to 60 GBytes and 36 Intel Xeon CPU cores for more than 500 nodes.

A Kubernetes cluster consists of the components that represent the control plane and a set of machines called nodes, sometimes referred to as worker nodes (Kubernetes Overview, n.d.). See Figure 8 - Kubernetes Components, taken from (Kubernetes Components, n.d.).

A pod is the smallest scheduling unit in Kubernetes and represents a set of containers that are tightly coupled, share resources and therefore run in the same worker node. Kubernetes runs workloads by assigning pods to nodes based on the resource usage and limits from the application. Each pod has its own IP address, and a default Kubernetes control plane runs its own DNS service for service to address resolution.





**Figure 8 - Kubernetes Components**

The Kubernetes control plane (master node) consists of:

- the database (etc. by default) that stores all cluster data,
- the API server (kube-api-server) that exposes the Kubernetes API and serves as its frontend,
- the scheduler (kube-scheduler) that watches for newly created pods and selects a node for them to run on,
- the controller processes runner (kube-controller-manager) that runs node, replication, endpoints and other controllers responsible for managing different elements of the cluster,
- the cloud-specific processes runner (cloud-controller manager) that runs processes specific to the cloud provider,
- for high availability deployments, the master node is configured on three separate machines.

The Kubernetes data plane (worker node) consists of:

- kubelet is the agent that accepts pod specifications from the control plane and runs them on the local container runtime (e.g. docker),
- kube-proxy is a network proxy that implements the Kubernetes Service concept, where one or more pods can sit behind a network service for load balancing purposes.

Kubernetes uses Network Plugins that run at the node level to configure container & pod network interfaces in the Linux OS and perform IP address management. By default, the kubelet is assigned with a plugin that supports a cluster-wide IP network. Container Network Interface (CNI) is a CNCF project that provides the specification and tools required to implement plugins to manage the allocation and deallocation of network resources for a container. Kubernetes supports plugins that adhere to the CNI specification and support can be extended as required by introducing new plugins for specific network functionality such as supporting container communications over VXLAN and MACVLAN (github.com/containernetworking, 2020).

Vanilla container-level networking is not suitable when orchestrating containerized data plane network functions (CNF) in a service chain or when more complex network overlays are required, such as creation and management of VXLAN tunnels that extend beyond a single Kubernetes cluster. Network Service Mesh is a cloud-native project that adds network capabilities to the Kubernetes ecosystem to enable dynamic cross-connections between local and remote CNFs (network service mesh, 2020). It offers an

API to establish connectivity between network services in an abstract way and provides policy-based service function chaining.

### 4.3.2. Kubernetes Distributions

Since a default Kubernetes installation (often referred to as “vanilla” Kubernetes) cannot be used without installing additional components, it is recommended to use one of the free or commercial Kubernetes distributions. A well-chosen Kubernetes distribution instead of the “vanilla” one reduces the operator’s dependency on Kubernetes experts and offloads a lot of installation and configuration work, this is definitely a preferable choice, at least until the operator develop in-house expertise and automation tools with Kubernetes. A broad categorization of distributions is illustrated in Figure 9 - Kubernetes Distribution Models.

There are several dozen Kubernetes distributions recognized by the (CNCF), and it is recommended to select from this list according to the operator’s requirements (CNCF Cloud Native Landscape, n.d.). We describe a couple of distributions below that are relevant for edge cloud.

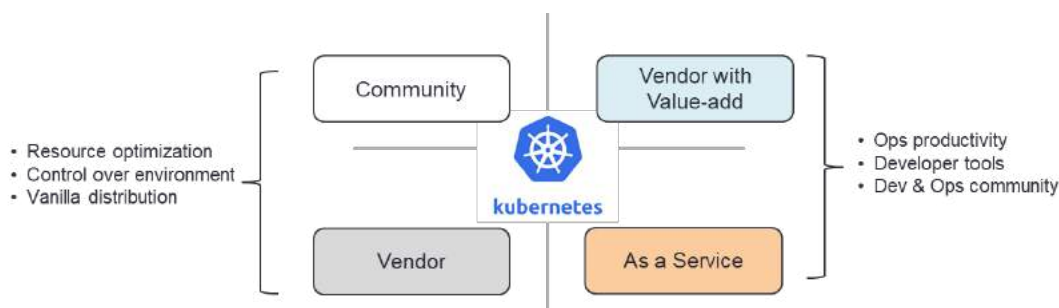


Figure 9 - Kubernetes Distribution Models

### 4.3.3. Vanilla Kubernetes Distribution

Kubernetes sources are available at GitHub (<https://github.com/kubernetes/kubernetes>) and many organizations take this source “as is”, build it for different platforms and distribute it without any significant addition. This form of distribution makes it more convenient than building it from the source and has a clear correlation between such distribution and a tagged Kubernetes version (see <https://github.com/kubernetes/kubernetes/tags>). Canonicals “Charmed Kubernetes” distribution (<https://jaas.ai/canonical-kubernetes>) is an example of a vanilla distribution. This distribution has broad applicability across Telco and MSO workloads without specialized compute requirements for containers.

### 4.3.4. Kubernetes with Value-add Capabilities

Some commercial organizations take the Kubernetes source and add significant functionality to it, making a commercial distribution with dedicated support. Kubernetes distributions such as Red Hat OpenShift (<https://www.openshift.com/>) and Rancher (<https://rancher.com/>) provide installers for fully automated cluster deployment as well as abstractions and tools for DevOps processes and CI/CD pipelines.

Both Red Hat OpenShift and Rancher use special, container-oriented Linux distributions for worker nodes (CoreOS and RancherOS respectively). Some of these distributions offer a Cluster API capability to manage multiple Kubernetes clusters that can be deployed on-prem, private cloud or public cloud.

#### **4.3.5. Kubernetes as a Service**

Kubernetes as a service is offered by all major cloud providers, AWS, Microsoft Azure, IBM and Google Cloud Platform. In this kind of distribution, the cloud service provider takes care of everything related to the Kubernetes version, allocation, and installation of master (control planes) and worker nodes, and all the underlying compute, storage and network resources. This type of managed Kubernetes service is managed from the cloud providing the ability to unify on-prem and cloud hosted clusters and flexibly deploy containers across a hybrid cloud.

This type of Kubernetes deployment makes practical sense for back-end management workloads that oversee an operator's network but are not suitable for infrastructure edge user plane workloads. In some cases, the cloud provider supports edge cloud environments that can be collocated in a hub/headend or operator edge data center to run enterprise services or network control plane functions at the metro level. When considering this service, one must clearly understand the operational benefits, managed services costs and the tradeoffs associated with lack of control and visibility of the underlying infrastructure.

#### **4.3.6. Lightweight Kubernetes for Edge Cloud and IoT**

A Kubernetes cluster consists of one or more master nodes and one or more worker nodes. Each node has Kubernetes binaries that dictate the role and behavior of the node. These binaries come with significant overhead, which makes them impractical for certain edge deployments where the underlying hardware is limited in compute and memory resources. Lightweight Kubernetes distributions address this class of hosting environment by stripping off functionality that is not required in small scale, single node use cases.

One of such distribution is k3s by Rancher built for IoT and edge compute (k3s.io, 2020). This distribution builds Kubernetes from a reduced source tree and changes its binaries structure with the single goal of making it as small as possible. It replaces the *etcd* database with reduced storage based on *sqlite3* and removes in-tree storage drivers and cloud providers. Also, it reduces the memory footprint by running many components inside a single process. This results in a “lightweight” distribution that is suitable to run both control and user planes on devices with limited resources, making it a great fit for Edge Cloud nodes to coincide with RPDs/RMDs.

Canonical makes another “lightweight” Kubernetes distribution called MicroK8s (MicroK8s, 2020). This distribution also targets IoT and edge computing. It makes most of Kubernetes options default, resulting in simplified installation, configuration, and updates.

#### **4.3.7. Kubernetes for Distributed Node Clusters**

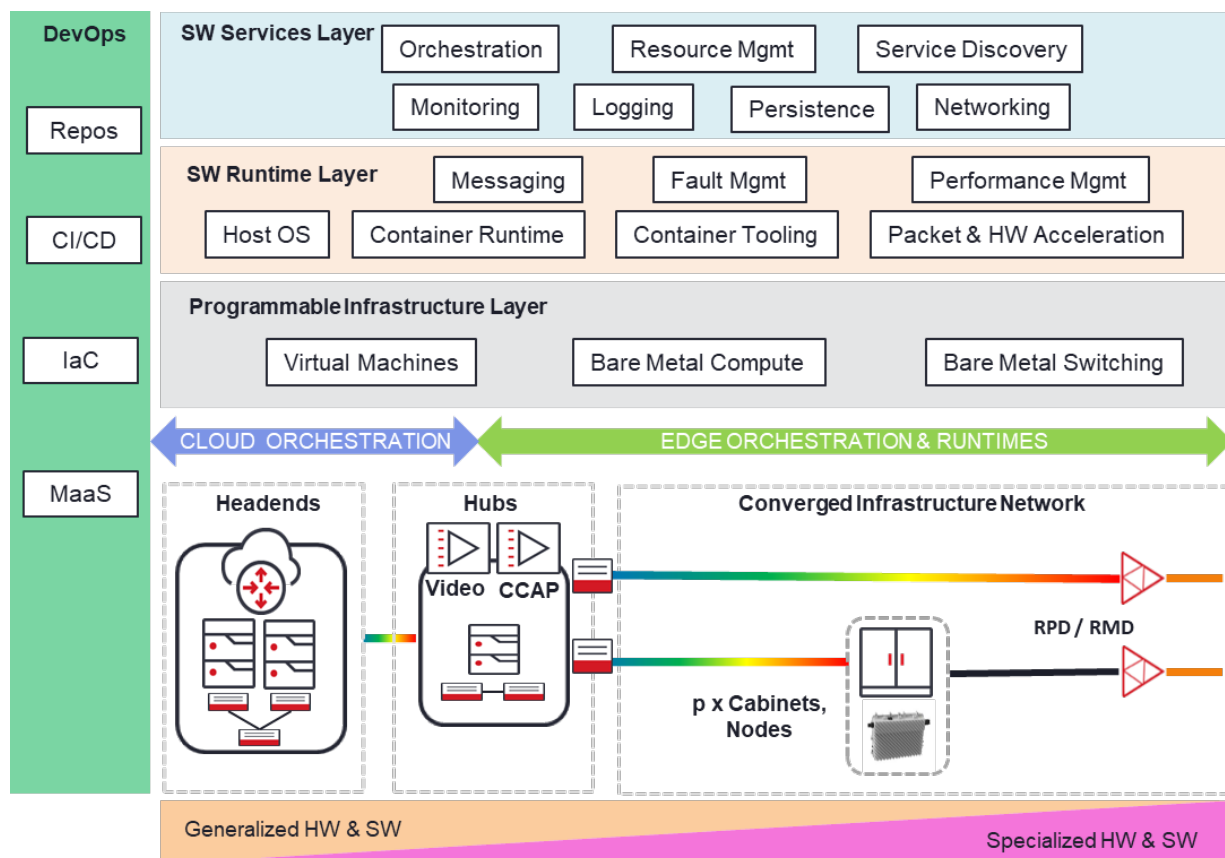
One more specialized Kubernetes distribution (or at least its approach) that should be considered. By default, Kubernetes requires a reliable connection between its master and worker nodes. KubeEdge is designed for environments where this connection may be intermittent and the Kubernetes master node (or control plane) is at the Cloud Tier while worker nodes are at the Edge Tier.

This distribution replaces the worker node's *kubelet* software with its own lightweight node agent called *edged*, and adds a special software layer on top of the master node. It changes how a worker node communicates with its master node located at the Cloud Tier, enabling it to work across non-reliable network conditions which the default Kubernetes control plane cannot tolerate.

Consider this alternative when there is a requirement for hosting the control plane in a central cloud location. The tradeoff for this is increased complexity and a significant deviation from the default Kubernetes, which locks the operator to this specialized distribution.

## 5. Building out a Distributed Edge Cloud

It is important to highlight the technical considerations in the deployment of virtualized infrastructure and network software applications such as in a DOCSIS DAA network. A proposed framework for this discussion is shown in Figure 10 - DAA Orchestration Framework.



**Figure 10 - DAA Orchestration Framework**

### 5.1. Cloud vs Edge Orchestration

As discussed before, Kubernetes orchestration of containers/pods has some underlying assumptions about homogeneity, network underlay and proximity of worker / compute nodes. The Kubernetes scheduler needs reliable and fast connectivity to its compute nodes, in this case across the access network to reach Hub and RPD/RMD nodes. Kubernetes networking also assumes that the underlying switching infrastructure support a flat IP domain interconnecting the physical host nodes.

There are primarily three design alternatives:

1. Define Kubernetes clusters according to homogeneous resource type and functional requirements: e.g. nodes with specialized resources to support RPD/RMD functions belong to the same cluster. This means multiple edge clouds consisting of either general purpose compute or specialized compute & accelerator are managed by different Kubernetes instances. This adds complexity to managing distinct clusters, but it allows for independent scaling of cluster and the interconnection network
2. Define Kubernetes clusters according to a flat physical topology of the interconnection network and including heterogeneous resources within the same cluster, i.e. general and special purpose resources. This simplifies the management of the centralized Kubernetes control plane itself, but complicates the scheduling algorithms needed to make container/pod placement decisions
3. Hybrid approach of the above two options to allow for incremental scaling and performance management of the network as services and subscriber density increases.

Similar issues arise with the default deployments of OpenStack because it was not designed for distributed edge clouds; resource heavy control planes need to be deployed at multiple locations with another management layer required to coordinate between different clouds. An end-to-end network service orchestration capability can be used to unify and stitch together services across multiple edge clouds.

## 5.2. Network Functions and Runtimes

One of the most important techniques for working with network functions at the edge is to treat them as individual building blocks or “microservices” with well-defined interfaces for configuration and user plane stitching. This provides operational flexibility to introduce new services and software updates, when and as needed, in an automated fashion.

Whether network function is control or user plane oriented, they will rely on standard communications protocols for interfacing with platform services or other network functions. Data transfers using RESTful mechanisms is the most widely used style implemented as JavaScript Object Notation (JSON) over HTTP/1.1. While this provides simplicity for clients, it does require a web server embedded in the network functions. Messaging using Remote Procedure Call (RPC) is protocol agnostic and offers direct communication between clients and server applications. Brokered messaging is another alternative usually used in high performance publish-subscriber communication models that requires a message broker. And lastly for monitoring applications the Message Queuing Telemetry Transport (MQTT) protocol is typically used. It is critical that VNF and CNF vendors align on standard protocols and interfaces to simplify integration and allow for optimization of messaging systems in the runtime environment.

Selection of the runtime environment for VNFs and CNFs is dependent on whether network functions are optimized for central cloud, edge cloud or specialized appliance environments. While these choices provide similar deployment agility from a Kubernetes scheduler, the runtime OS, guaranteed priority scheduling, and HW acceleration abstraction capabilities are critical for real-time sensitive network functions. Servers with accelerators (GPUs, FPGAs) or specialized ASIC-based appliances with generalized compute need a high reliability runtime to guarantee network and CPU resources, and provide a high-performance data channels between VNF/CNFs. RPD and RMD nodes with limited resources can use lightweight Kubernetes distributions and a specialized runtime for access to packet processing, HW acceleration features and policy-driven CNF chaining.

### 5.3. Automated Operations & Business Network Intent

Manual deployment and operation of edge cloud infrastructure is not practical, especially with the scale of MSO networks, services and subscribers. Software delivery, VNF/CNF instantiation and replacement must be fully automated by a software services, orchestration, and monitoring layer. This is one of the reasons Kubernetes adoption has risen quickly; changes in the infrastructure, application scale up/down and recovery on failures are fully automated based on application intent provided through a manifest.

Modern automation tools can be classified as supporting declarative or imperative style of programming. With declarative style, the infrastructure is described as “what” needs to be built, vs the imperative – “how”. In order to support immutable infrastructure concepts described earlier, a declarative “what” style of tools (like Terraform) that support the notion of “Infrastructure As Code” should be used for managing infrastructure. This approach decouples infrastructure description from the tools used to build it, making it possible to automate across different compute and network systems. This is beneficial to maintain a common baseline of infrastructure software across a heterogeneous environment of Headend/Hub and remote RPD/RMD nodes.

A network fabric supporting edge cloud needs a holistic, data-driven closed loop approach that automates key business processes that span IT systems like planning, fulfillment and assurance and network lifecycle operations such system connectivity & network functions configuration, scaling and healing. Describing network intent in the form of operational state allows network algorithms to determine whether or not the network is deviating from its intended state and therefore take corrective actions, such as proactively re-routing traffic to avoid congestion or alerting the operator of impending network outages before they occur. Kubernetes intent-driven orchestration of containers is well suited to support this model as long as the scheduling of end-user applications can influence the allocation of the underlying network infrastructure, whether physical or virtualized. This enables optimal use of network resources based on application needs from core cloud to edge cloud to access.

One key area that requires special extensions to Kubernetes is the scheduling & placement of containerized functions based on not only CPU, memory, and specialized HW availability but also networking constraints that exist in a distributed edge cloud environment across headend/hub and remote nodes. The basic Kubernetes scheduler assesses in real-time CPU and memory resource requests and limits to decide on resource allocation and contention management. This works well in data center applications where the cluster fabric is overprovisioned and fully meshed, so placement decisions can be made virtually ignoring available bandwidth or effective round-trip-time between nodes. A network aware scheduler is required to maintain not only container resource usage and limits for compute and memory, but also topological constraints, specialized accelerators, network capacity, packet loss, latency and jitter metrics.

### 5.4. Application Repositories

Each application and virtual network function need to be placed in one or more repositories where the orchestration layer can retrieve and deploy across the infrastructure. This repository provides a controlled place where version-controlled artifacts can be maintained. This capability is vital for mitigation of issues that arise after a botched deployment. Automation must be able to restore a system to the last working state without network operator intervention. It is highly desirable to have this repository integrated with a software development Continuous Integration/Continuous Delivery (CI/CD) process or vendor pipelines to accelerate tactical deployment of new features and fixes.

Another proven practice in automated software operations is the blue-green model of deployment, where two identical software stacks are maintained, with the first (blue) being alive and the second (green) being on standby. Switching between stacks should be automated and without any downtime to the end-user or network service. When new components are released, they are added only to the green stack which becomes operational, while the blue remains on standby. Once the green stack has demonstrated no production issues, the blue stack is updated. However, if the green stack fails, then “old” blue stack becomes operational, while the green one is taken offline for troubleshooting.

## 5.5. Undercloud Architecture

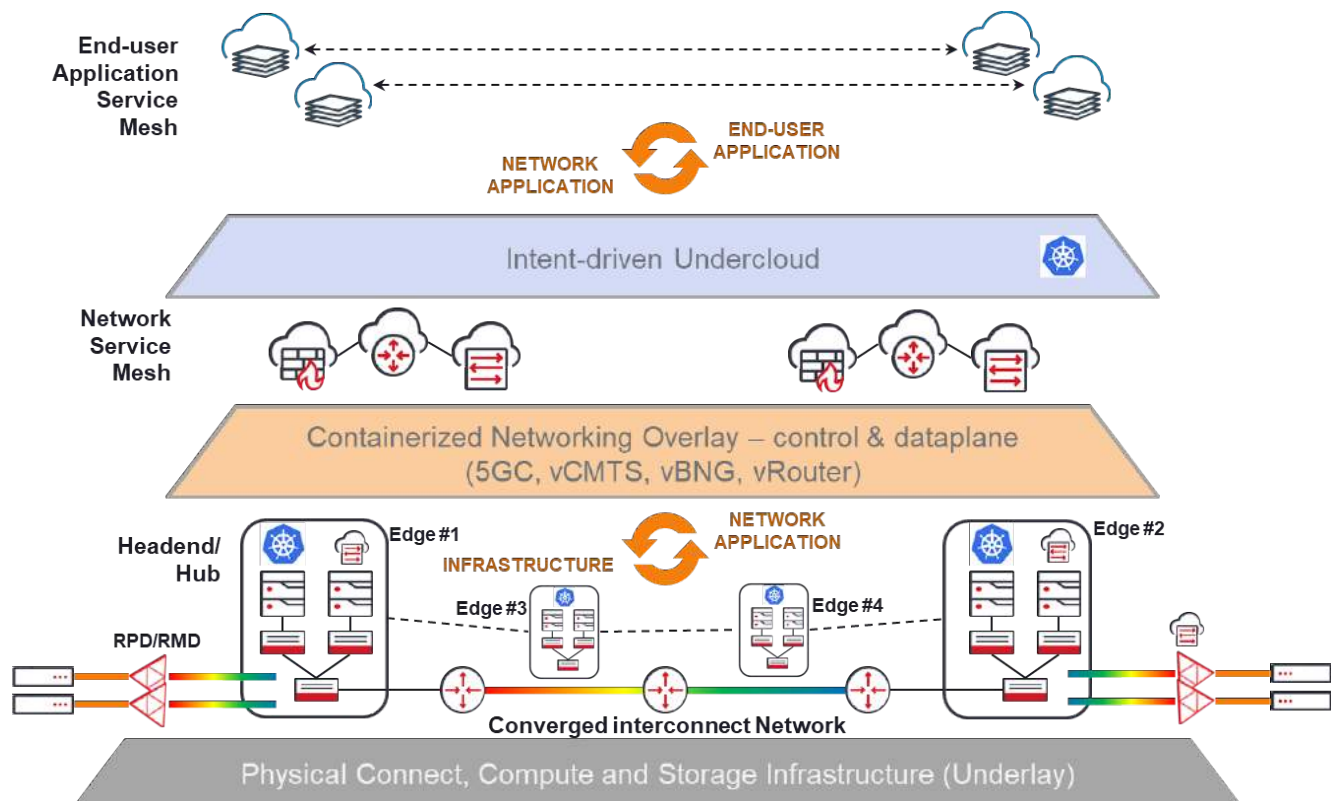
In advanced software-driven networks, there is an expectation that the end-user applications riding over the network can influence network behavior and resource allocation to meet QoS, security and reliability requirements. The dynamic nature of containerized multi-cloud applications necessitates a flexible and intent-driven approach to automatically allocate, configure, monitor and scale compute and network connectivity resources between the physical infrastructure underlay and the network functions overlay as depicted in Figure 11 – Intent-Driven Undercloud.

The Physical Infrastructure Underlay (PIU) is made up of all the physical resources installed at headend/hub and remote node locations for the transmission, switching and processing of video and data services delivered to subscribers. This physical underlay is SDN-controlled and highly instrumented to enable high fidelity telemetry to be used by the Containerized Network Overlay (CNO) to ensure network applications and functions are meeting the demands placed on them by the End-User Application Service Mesh (ASM).

The CNO is made up of CNFs, and VNFs where appropriate, to deliver network control and user plane services such as vBNG, 5GC and UPF, vRouting and vCMTS. Due to the multi-layer nature of the Converged Interconnect Network (CIN), a Network Service Mesh (NSM) associated with the networking Kubernetes clusters can be used to drive policy-based service chaining and configuration of the CNO and PIU layers. These service chains, both control and user-plane centric, support specific services that are specified as “network intent” by the Application Service Mesh. This technique ensures that the CNO layer is making closed loop decisions with dynamic network information to operate within the policies specified by the operator while maintaining the intent of the network services specified by the ASM. This coordination of intents is based on the use of application manifests that specify resource requirements and constraints to the Kubernetes control plane.

It should be noted that the CNO and ASM layers are orchestrated and managed by different Kubernetes control planes which may be operated by the network operator or a separate application/cloud provider with integration into the operator’s Intent-Driven Undercloud (IDU) and direct network peering into the Headend/Hub sites.

The Intent-Driven Undercloud (IDU) is a policy layer that maps application networking intent into a Network Service Mesh (NSM) intent that oversees multi-domain, multi-vendor, multi-layer functions such as available inventory and capacity of resources, resource management, orchestration and monitoring of end-to-end connectivity services. The IDU initiates real-time deployment & configuration, automatically allocating physical compute, store and network resources and stitching an end-to-end connectivity service that supports the CNO layer.



**Figure 11 – Intent-Driven Undercloud**

### 5.5.1. Undercloud Control Plane


When orchestrating the undercloud with Kubernetes, each cluster is part of a headend/hub data center, where the master nodes (or control plane) reside. This control plane site uses the underlay CIN network for reliable and fast connections to the remote worker nodes or RPD/RMD nodes. This centralized architecture enables the master nodes to be configured for high availability and to scale where compute resources are homogeneous and plentiful.

Two approaches can be used to cluster compute nodes in the DAA network:

1. Organize clusters by resource type constraints
2. Organize clusters by physical connectivity constraints

The first approach, as illustrated in Figure 12 - Homogeneous Resource Scheduler, takes advantage of native Kubernetes scheduling features for deploying CNFs within the same pool of resources allowing an external end-to-end service orchestration component to abstract the physical topology dependencies in the access network and request CNF deployments according to network centric constraints such as bandwidth availability, latency, packet loss, and jitter. This method allows for independent scaling & replacement of cluster nodes, clear separation of concerns between localized compute functions and decentralized network connectivity and enables independent scaling of the CIN topology. The end-to-end Service Orchestration component is responsible for calculating network paths and ensuring application intent is met through the interconnection of physical and virtual/container functions across clusters.






**Figure 12 - Homogeneous Resource Scheduler**

The second approach, as illustrated in Figure 13 - Heterogeneous Resource Scheduler, requires that Kubernetes be aware of network topology to make pod scheduling decisions. As discussed before, when Kubernetes is used for orchestrating the CNO across a heterogeneous network pool, its scheduler assigns pods to remote nodes based on CPU and memory resources only, potentially making sub-optimal decisions due to dynamic network latency and performance characteristics that it is not aware of.

To overcome this problem, the Kubernetes scheduler can be extended with capabilities to make network-centric decisions (Kubernetes Scheduler Extensions, 2020). This network-aware scheduler can use labeling mechanisms across the physical underlay to build a view of network topology, latency & BW utilization underpinning worker nodes. Additionally, it can label specialized resources such as FPGAs, GPUs and smart NICs to filter and select nodes given the CNF resource requirements.



**Figure 13 - Heterogeneous Resource Scheduler**

This label-based, custom scheduler extension mechanism can be used to place CNFs to highly optimized nodes with run-time operating systems, available specialized software as Data Plane Development Kit (DPDK), Single-root input/output virtualization SR-IOV and various accelerators.

### **5.5.2. Undercloud Resources & Edge Optimized Runtime**

The resource-constrained reality of the distributed edge cloud nodes means careful consideration must be given to CPU, memory, and storage requirements for Kubernetes. Since centralization of the control plane is possible, the remaining concern is with the runtime environment footprint and Kubernetes agents to coordinate with the centralized master.

We have discussed several options in this paper, including lightweight Kubernetes distributions such as k3s, MicroK8s and KubeEdge to address Kubernetes agents running on worker nodes. The other challenge is addressing the container runtime itself, as discussed in Virtual and Cloud Native Network Functions. This runtime package should be based on a real-time Linux distribution built using the Yocto Project (Yocto Project, 2020).

As mentioned above, highly optimized resource nodes with real-time operating systems and specialized software and hardware can be included into pod scheduling decisions using Kubernetes labels and custom scheduler extensions.

Within CableLabs, there is a new initiative, called Project Adrenaline, which is harnessing momentum to address the management of heterogeneous accelerators available at different locations in the cable access network. This project aims to promote technologies and architectures that enable a distributed & heterogeneous edge compute fabric to support dynamic placement of workloads (Levensalor & Stuart, 2020). The ability to orchestrate workloads and abstract the use of accelerator resources through an edge optimized Kubernetes runtime is extremely beneficial for the application developer community; this initiative will accelerate application design cycles and deployment of new features and bug fixes, independently of the underlying infrastructure allowing for concurrent innovation and cloud-style delivery.

## **5.6. Open-Source Building Blocks**

In previous sections, we have described the use of certain open source software components to build edge cloud infrastructure. In this section, we will provide a high-level overview of relevant open source projects and pointers for additional information. Open source software greatly reduces the price for solutions and components when the same codebase is used by several businesses, and these businesses coordinate the development effort and prioritization of features. Several organizations provide a means for coordination of development of open source software, usually in the form of a membership. The Linux Foundation is one of the largest and best known open source organizations and it includes several “suborganizations” with focus on specific areas of technology, one of which is edge.

### **5.6.1. LF Edge**

The Linux Foundation *Edge* was announced in Jan 2019 as an “umbrella organization to establish an open, interoperable framework for edge computing independent of hardware, silicon, cloud, or operating system” and includes over 30 “Premier” members from the operator, cloud and vendor community and over 40 “General” and “Associate” members (Linux foundation edge, 2020).

One of the most relevant for cable MSOs is the Akraino Edge Stack project.

### 5.6.1.1. Akraino Edge Stack

The Akraino Edge Stack intends to develop a fully integrated edge infrastructure solution for the Edge tier (LFEdge Akraino, 2020). This project consists of two main elements:

1. Blueprints as declarative configurations of entire software stacks to address specific use cases,
2. Software for common components declared in blueprints

The Akraino edge stack is an open-source software stack that improves the state of edge cloud infrastructure for operators, service providers and IoT networks. This edge stack can be viewed as the runtime and infrastructure layers for VNFs and CNFs. Akraino Blueprints are divided into two groups: approved and proposals, that are structured into several families. Each blueprint targets a very specific use case and a very specific deployment size, referred to as “Point of Delivery” (Akraino PODs, 2020).

Each blueprint uses the Akraino reference software stack illustrated in Figure 14 - Akraino Software Stack (The New Intelligent Edge - Akraino Edge Stack Overview, 2018).

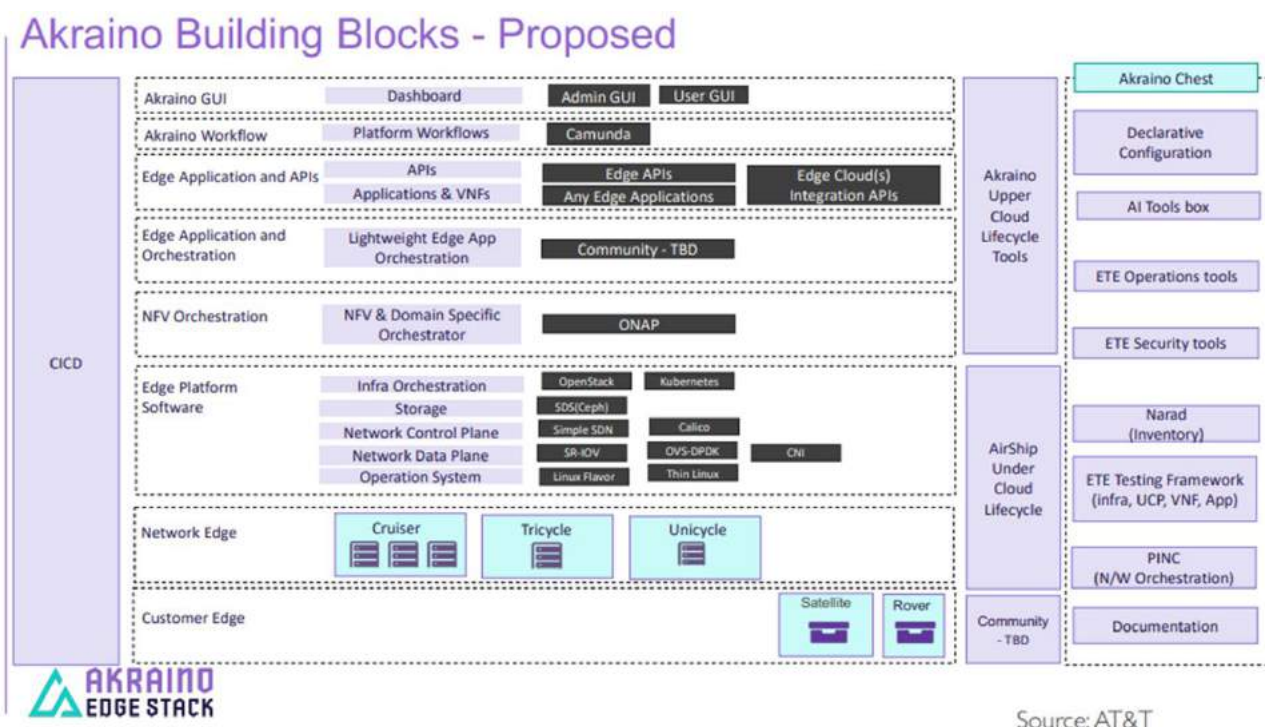


Figure 14 - Akraino Software Stack

Akraino uses an edge cloud architecture model with the full control plane located at the infrastructure edge. A blueprint represents a standard model of deployment for various operator sites: central regional and edge. While the blueprints cover a wide range of use cases with Multi-access Edge Computing (MEC) and 5G vRAN (Virtualized Radio Access Network) applicable to cable MSOs, there are no blueprints to address different control plane deployments.

A relevant blueprint for this paper is the Kubernetes-Native Infrastructure (KNI) Blueprint Family. The KNI is optimized for Kubernetes-native workloads and also allows hybrid deployments (CNF & VNF) using KubeVirt, a technology that allows VMs to run as a pod inside a Kubernetes cluster. There are currently two KNI blueprints in progress:

- Provider Access Edge (PAE) optimized for real-time and high performance vRAN and MEC workloads
- Industrial Edge (IE) optimized for small footprint and low latency for IoT, serverless and machine learning workloads

It is noteworthy that KNI uses a commercial Kubernetes distribution called “Red Hat OpenShift” (<https://www.openshift.com/>) and the Cluster API (<https://cluster-api.sigs.k8s.io/>) to deploy a Kubernetes cluster. The Cluster API is declarative and uses tooling to simplify provisioning, upgrading, and operating multiple Kubernetes clusters. This makes it much more suitable for the distributed edge computing deployments.

### **5.6.2. Cloud Native Computing Foundation**

The Cloud Native Computing Foundation (CNCF) hosts open-source software components for cloud native applications (<https://www.cncf.io/>). CNCF hosts Kubernetes, however there are over 1,400 projects, product, or technologies under the CNCF umbrella, see (CNCF Landscape). It is recommended to get familiar with categories of technologies and products and to see how CNCF suggests combining them into a solution.

## **5.7. Commercial Cloud Platforms for Edge Computing**

In its simplest form integration with a cloud platform is a matter of accessing specific endpoints on the internet, but most of cloud platform providers offer more tightly coupled software supporting computations at the Edge Tier.

When looking at this software it helps to understand that the Internet of Things (IoT) was the very first use case for edge computing that was offered by commercial and open-source cloud platforms. The IoT use case requires very specific cloud-based components (like a thing registry) and messaging protocol (MQTT), and Edge Tier nodes usually serve as IoT gateways. Also, this use case requires data processing as close as possible to the source. Such processing can be done as a dedicated process at the edge node, in the form of a standalone application, a container, a virtual machine or a serverless application. Serverless application frameworks are increasing in popularity because they do not require packaging, easily fit into event-driven design and use compute resources only when active.

### **5.7.1. AWS IoT Greengrass**

The AWS IoT Greengrass consists of a binary that is installed on a node (a Linux server, for instance) at the Edge Tier. After installation and registering with the AWS IoT cloud service, it provides MQTT messaging service to IoT devices. Messages can be processed right at the node by a serverless application (integrated with AWS Lambda service) or forwarded to the cloud for consumption by other AWS services. AWS IoT Greengrass also integrates with artificial intelligence software, providing means to run models pre-trained at AWS or elsewhere.

### **5.7.2. AWS Outpost**

AWS Outpost is a ready-to-use edge cloud infrastructure. It provides several compute, storage and networking services that are enough to run an Edge Tier based datacenter. Typically, an Outpost rack of server is deployed on the enterprise premise or the operator's network to peer directly with access networks such as 5G and perform real-time edge processing functions. Outpost is offered as a managed service and is considered an extension of AWS cloud regions, making it possible to seamlessly deploy applications across core and edge clouds.

### **5.7.3. Azure IoT Edge**

Microsoft Azure IoT Edge is a binary that extend the Azure IoT Hub to the edge. As AWS IoT Greengrass, it can serve as IoT gateway, and can run containers and artificial intelligence software, all integrated with respective Azure services.

### **5.7.4. Azure Stack Hub**

Microsoft Azure Stack Hub is another ready-to-use edge cloud offer, very similar to AWS Outpost, except that it is only a software stack that run on commodity servers. It provides several compute, storage and networking services that are enough to offer IaaS and PaaS services at edge locations or local zones.

### **5.7.5. Google Cloud Anthos**

Google offers a different approach to enable edge and multi-cloud environments. Google Anthos is a control plane and run-time Kubernetes environment that unifies delivery of containerized applications across a wide variety of public (e.g. AWS) and private cloud (e.g. VMware) environments.

## **6. Conclusion**

Cable operators' investments to modernize access networks and move towards DAA is opening up new opportunities to transform their operations models and differentiate their network services with support for edge-centric Enterprise applications. Embracing cloud-native principles and an applications-first mindset is critical to the success of this transformation while simultaneously creating new revenue streams for Edge Cloud applications. Adapting Kubernetes orchestration and containerization of network and application functions are foundational first steps, which when coupled with an Intent-Driven Undercloud as defined in this paper, creates an adaptive and application aware network built for dynamic scale, business agility, operational efficiency and service innovation. This strategy enables low-latency and high bandwidth on-ramps to edge cloud resources at Hubsites and DAA locations where dynamic application demands can be satisfied through intelligent placement of network & application topologies.

## **Abbreviations**

5G	Fifth Generation cellular network technology
AMQP	Advanced Message Queuing Protocol
API	Application Programming Interface
AWS	Amazon Web Services
BSS	Business Support System

CIN	Converged Interconnect Network
CMTS	Cable Modem Termination System
CNCF	Cloud Native Computing Foundation
CNF	Cloud-native Network Function
CNI	Container Network Interface
CNO	Containerized Network Overlay
cRAN	Centralized Radio Access Network
CRI	Container Runtime Interface
DEC	Distributed Edge Computing
DPDK	Data Plane Development Kit
ENF	Edge native Network Function
GCP	Google Cloud Platform
gRPC	gRPC Remote Procedure Calls
HTTP	Hypertext Transfer Protocol
IT	Information technology
IoT	Internet of Things
IP	Internet Protocol
JSON	JavaScript Object Notation
K8s	Kubernetes
KNI	Kubernetes-Native Infrastructure
KVM	Kernel-based Virtual Machine
LAN	Local Area Network
LF	Linux Foundation
LLD	Low Latency DOCSIS
MAN	Metropolitan Area Network
MEC	Multi-access Edge Computing
MQTT	Message Queuing Telemetry Transport
MSO	Multiple-System Operator
NIC	Network Interface Card
NFV	Network Function Virtualization
NSM	Network Service Mesh
OCI	Open Container Initiative
ONAP	Open Network Automation Platform
OPNFV	Open Platform for NFV
OVN	Open Virtual Network
OSS	Operations Support System
PIU	Physical Infrastructure Underlay
POD	Point of Delivery
REST	Representational State Transfer
RMD	Remote MACPHY Device
RPC	Remote Procedure Call
RPD	Remote PHY Device

VNF	Virtualized Network Function
vBNG	Virtual Broadband Network Gateway
vFW	Virtual Firewall
vRAN	Virtualized Radio Access Network
WAN	Wide Area Network
XML	Extensible Markup Language
YAML	YAML Ain't Markup Language

## References

*Akraino PODs*. (2020). Retrieved from <https://wiki.akraino.org/pages/viewpage.action?pageId=1147248>

Alusha, D. (2019). *Cloud-native computing in 5G networks*. Oyster Bay, NY: ABI research for visionaries.

*CNCF Cloud Native Landscape*. (n.d.). Retrieved from CNCF:  
<https://landscape.cncf.io/category=certified-kubernetes-distribution&format=card-mode&grouping=category>

*CNCF Landscape*. (n.d.). Retrieved from Cloud Native Computing Foundation: <https://landscape.cncf.io>

*coreos.com/rkt/*. (n.d.). Retrieved from [coreos.com/rkt/](https://coreos.com/rkt/): <https://coreos.com/rkt/>

*Evolution to Distributed Access Architectures*. (n.d.). Retrieved from COMMScope:  
<https://www.commscope.com/solutions/fixed-access-networks/distributed-access-architecture/>

*github.com/container networking*. (2020). Retrieved from <https://github.com/container networking/cni>

*github.com/opencontainers/runc*. (n.d.). Retrieved from <https://github.com/opencontainers/runc>

*k3s.io*. (2020). Retrieved from [k3s.io](https://k3s.io/): <https://k3s.io/>

*Kubernest Best Praactices*. (n.d.). Retrieved from [Kubernest.io](https://kubernetes.io/docs/setup/best-practices/cluster-large/#size-of-master-and-master-components): <https://kubernetes.io/docs/setup/best-practices/cluster-large/#size-of-master-and-master-components>

*Kubernetes Components*. (n.d.). Retrieved from [Kubernetes.io](https://kubernetes.io/docs/concepts/overview/components/):  
<https://kubernetes.io/docs/concepts/overview/components/>

*Kubernetes Overview*. (n.d.). Retrieved from [Kubernetes.io](https://kubernetes.io/docs/concepts/overview/what-is-kubernetes/):  
<https://kubernetes.io/docs/concepts/overview/what-is-kubernetes/>

*Kubernetes Scheduler Extensions*. (2020). Retrieved from <https://kubernetes.io/docs/concepts/extend-kubernetes/#scheduler-extensions>

*Kubernetes, automated container deployment, scaling, and management*. (2020). Retrieved from [Kubernetes, automated container deployment, scaling, and management](https://kubernetes.io/): <https://kubernetes.io/>

Levensalor, R., & Stuart, C. (2020, July). *The Modular, Virtualized Edge for the Cable Access Network*. Retrieved from Adrenaline™ Project: <https://openadrenaline.com/>

- LF Edge. (2020, June 20). *Open Glossary of Edge Computing 2.1.0*. Retrieved from LF Edge: <https://github.com/State-of-the-Edge/glossary/blob/master/edge-glossary.md>
- LFEdge Akraino. (2020). Retrieved from lfedge.org: <https://www.lfedge.org/projects/akraino>
- Linux foundation edge. (2020). Retrieved from Linux foundation edge: <https://www.lfedge.org>
- MicroK8s. (2020). Retrieved from MicroK8s: <https://microk8s.io/>
- Namiot, D., & Sneps-Sneppé, M. (2014). On -Micro-services Architecture. *International Journal of Open Information Technologies*, 24-27.
- network service mesh. (2020). Retrieved from network service mesh: <https://networkservicemesh.io/>
- The Converged Interconnect Network. (2020). Retrieved from <https://www.ciena.com/insights/white-papers/the-converged-interconnect-network.html?aliId=eyJpIjoiMGFuUG9ZeTFwV2djdR0TyIsInQiOiJsb3hSd05oZnZ4d2V5bGVZTXlnVG9BPT0ifQ%253D%253D>
- The New Intelligent Edge - Akraino Edge Stack Overview. (2018). Retrieved from <https://object-storage-ca-ymq-1.vexxhost.net/swift/v1/6e4619c416ff4bd19e1c087f27a43eea/www-assets-prod/summits/24/presentations/21275/slides/Akraino-OverviewOpenStackv2.pdf>
- Whie, G., Sundaresan, K., & Briscoe, B. (2019). *Low Latency DOCSIS: Technology Overview*. Retrieved from <https://www.cablelabs.com/technologies/low-latency-docsis>
- Yocto Project. (2020). Retrieved from <https://www.yoctoproject.org/>



AUTHOR INDEX  
2020 Technical Paper Proceedings

Abboud, Claude Bou .....	1601	Carro, Gabriel .....	1949
Abu-Hijleh, Omar .....	1709	Carroll, Chuck .....	729
Alapati, Venkata Somi.....	1342	Chan, Tat .....	537
Al-Banna, Ayham .....	81	Chang, Gee-Kung .....	484
Anderson, Aaron.....	395	Chapman, John T.....	827, 1212, 1618, 1739
Anderson, Rob .....	1089	Check, William A.....	2231
Andréoli-Fang, Jennifer.....	1739	Cheevers, Charles .....	290, 344, 1475
Archambault, Sylvain .....	1739	Chen, You-Wei.....	484
Ariesen, Jan .....	1370	Chrostowski, John .....	788, 1233, 2068
Bajwa, Anjan .....	562	Cloonan, Ruth.....	1370
Barker Jr., Bruce E. ....	1601	Cloonan, Tom.....	81, 518, 1370, 1709, 1979
Baruch, Eli.....	2309	Combs, Doug.....	941
Beaudin, Adrian.....	164	Combs, Jason .....	2046
Beesley, Bill .....	1461, 1469	Cornaglia, Bruno .....	1739
Bencheikh, Ahmed .....	1475, 1739	Cruickshank, Robert .....	12, 1073
Bestermann, Jay.....	2025	Curran, Anthony .....	811
Birnbaum, Jack .....	846	Daoud, Mohamed .....	1827
Blaser, Paul.....	1739	Davis, Drew .....	484, 1739
Bourg, Kevin .....	2331	Day, Chris.....	750
Bracker, Will .....	780	Dearborn, Colin .....	605
Briard, Xavier .....	2291	Dharanikota, Ayarah.....	344, 383
Brzozowski, John Jason.....	1881	Dharanikota, Sudheer .....	344, 383
Campbell, Chad .....	198	Dhillon, Parmjit .....	1827
Campbell, Ian .....	1739	DiGiacomo, Derek.....	12
Capuano, Simone.....	966	DiMicelli, Tom .....	902

Dokter, Mark .....	172	Goeringer, Steven .....	770, 780
Dolan, John.....	1271	Gohman, Greg .....	2008
Donofrio, Louis .....	2037	Gonsalves, Robert.....	846, 966
Downey, John J. ....	316, 331	Gowans, Paul.....	871
Eccles, Ryan .....	1418	Grayson, Mark .....	1739
Epstein, Steven .....	1	Guibene, Wael .....	1525, 1590
Fame, Hany.....	2172	Haefner, Kyle .....	1197
Fedorov, Dmitri .....	2342	Hallberg, Jacob .....	2200
Fernandes, João Pedro .....	1144	Harb, Maher.....	244, 1654
Ferreira, Jude .....	244, 1654	Hayes, Keith R.....	1448
Finkelstein, Jeff .....	484, 518	Heaton, Eric.....	494
Fiorenzo, Mariela.....	1949	Hewavithana, Thushara .....	674
Fish, Roger .....	198	Hmimy, Hossam .....	1525, 1590, 1697
Flesch, J.R. ....	290	Hohman, Kyle.....	592
Florenz, Ken .....	2025	Hoole, Elliott .....	2133
Foroughi, Nader.....	624, 674	Howald, Robert.....	1370, 1535, 2068
Francisco, Mark .....	2057	Hranac, Ron.....	198
Frederick, Andrew .....	2081, 2095	Hurley, Thomas .....	1271
Ganbar, Jambi.....	149	Hurtado, Omar .....	1949
Garcia, Albert .....	2222	Jindal, Manish.....	1697
Gaydos, Bob .....	918, 941, 966	Jones, Doug .....	518
Geary, David.....	12	Kakinada, Umamaheswar Achari .....	1697
Ger, Javier.....	1116	Kercher, Shawn .....	2200
Giladi, Alexander.....	1726	Khalilian, Michael .....	770
Gilbert, Ken .....	12	Kidani, Yoshitaka .....	891
Gladish, Jacob.....	1568	Kipp, Neill .....	1049
Glaser, Mike .....	12, 1271	Klatsky, Carl.....	788
Glennon, Steve .....	1172	Kolze, Tom.....	198

Koshy, Kamal.....	1475, 2291	Medvinsky, Alexander.....	537
Krauss, Simon.....	780	Medvinsky, Sasha.....	395
Kreisel, Michael.....	846	Menu, Eric.....	1739
Krishna, Karthik.....	149	Metts, Nicolas.....	1073
Kristoffersen, Even.....	198	Morley, Dave.....	1739
Kurkowski, Stuart.....	1049	Murphy, Arnold.....	1271
Lagacé, David.....	1739	Musat, Todd.....	729
Lalam, Massinissa.....	2291	Mutalik, Venk.....	918, 941, 966
Latini, Patricio Sebastian.....	704	Naveda, Marco.....	2342
Lee, Brian.....	484	Neeld, Erik.....	1601
Lewandowski, Benny.....	1233	Nickel, Ken.....	1271
Li, Xiaohua.....	1292	O'Keeffe, Frank.....	1909
Liu, Tong.....	827	Ovadia, Shlomo.....	63
Lu, Zhen.....	2172	Ozer, Sebnem.....	788, 1370
Lumbatis, Kurt.....	290, 1475	Page, Jason.....	471
Lund, Robert M.....	2108	Pala, Massimiliano.....	104
Lynch, Dan.....	63	Panciera, Eduardo.....	1739
Maenpaa, Timothy.....	420	Partridge, Jim.....	2231
Maricevic, Zoran.....	1909	Pasion, Jason.....	537
Markovic, Chris.....	1709	Patterson, Brian.....	12
Martens, Corwin.....	562	Paudel, Indira.....	2309
Martushev, Andy.....	811	Pavlich, Bryan.....	290
Matatyaou, Asaf.....	1330	Perron, Philippe.....	1739
Mathur, Tushar.....	1709, 2008	Pickering, Ladan.....	462
Matsumoto, Shuichi.....	891	Pinckernell, Nick.....	267
McFarland, William.....	990	Poltz, Damian.....	1739
McIntyre, Robert.....	446	Pratt, Craig.....	1568
Medlock, James.....	198	Prodan, Richard S.....	1845

Qiu, Xin .....	537	Seiden, Joshua .....	2192
Rakowsky, Lew .....	729	Shakil, Kashif .....	1342
Ramamurthy, Vignesh .....	2037	Shamsaasef, Rafie.....	395
Ramaswamy, Srinath V. ....	1644	Shen, Shuyi.....	484
Ranganathan, Raghu.....	2342	Shukla, Nishesh .....	2192
Ranganathan, Ram.....	1709, 2008	Sigman, Steve .....	1669
Rawat, Deependra.....	63	Singh, Pardeep .....	2172
Reale, David .....	1029	Snyder, Curtis .....	1073
Rebelo, Marcus.....	1899	Soeberg, Aleksander.....	198
Reyes, Elías Chavarría.....	1212	Spanbauer, Rick.....	966
Rice, Dan .....	244, 788, 918, 941, 966, 1233, ..... 1370, 1654, 2068	Spee, Rene .....	729
Righetti, Claudio.....	1949	Srivastava, Praveen.....	1739
Rome, Scott .....	267	Strobel, Rainer .....	674
Rose, Joshua .....	750	Subramanya, Karthik .....	1654
Rubinovitz, Rob.....	2231	Sundaresan, Karthik.....	1144, 1172
Rupe, Jason.....	35, 198, 624	Syed, Haider .....	2254
Saes, Claudio .....	1116	Szymanik, Colleen.....	2213
Saksena, Vikram.....	1418	Taylor, Kevin.....	770
Salinger, Jorge .....	918, 1669	Teague, John.....	12, 1271
Sanders, Kathryn .....	2108	Ten, Sergey.....	2331
Sandoval, Frank .....	1015	Thakore, Darshak.....	1568
Santangelo, Bryan.....	244, 1654	Thompson, Robert .....	1292, 1370, 2068
Sarawat, Vikas.....	1739	Tooley, Matthew.....	2231
Schauer, Paul .....	198, 1073, 2160	Tresness, Greg .....	1233
Schmitt, Matthew.....	1427	Trouillard, Géraldine .....	1739
Schwechel, Craig.....	1342	Ulm, John .....	1370, 1709, 1909, 1979
Seeger, Chris.....	127	Van Nice, Bruce .....	164, 172
		Ventriglia, Gary .....	846

Villarruel, Fernando X.....	1029
Vishnyakova, Anastasia.....	846
Vladyka, Andrii .....	1330
Vugumudi, Rohini .....	2172
Weerasinghe, Srilal.....	878
Weiss, Wesley .....	562
Welsko, Mark .....	12
Wheelock, Ian.....	344, 1475
White, Greg .....	1172
Wigley, Peter .....	2331
Wike, Pat .....	941
Williams, Tom.....	198
Winter, Eric .....	770
Woginrich, Tom .....	434

Wolcott, Larry .....	198, 846, 1370
Wolff, Peter .....	186
Wright, Ethan .....	1513
Yamany, Sameh.....	871
Yamashita, Hiroyuki.....	891
Yavuz, Mehmet .....	1891
Young, Gavin .....	1739
Zang, Qin.....	2037
Zeng, Helen .....	446
Zhang, Bob .....	2008
Zhang, Hongbiao .....	186
Zhou, Lei .....	2068
Zhou, Qi.....	484
Zhu, Jay .....	1144
Zhu, Jingjie .....	35

